SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2019

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2019

Herausgeber	Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Germany Telefon: +49 681 302 - 4396 Fax: +49 681 302 - 4397 E-Mail: service@dagstuhl.de
Registernummer	Amtsgericht Saarbrücken HRB 63800
Vorsitzender des Aufsichtsrates	Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen
Geschäftsführung	Prof. Raimund Seidel, Ph. D. Heike Meißner
Gesellschafter	Gesellschaft für Informatik e. V., Deutschland Universität des Saarlandes, Deutschland Technische Universität Kaiserslautern, Deutschland Karlsruher Institut für Technologie (KIT), Deutschland Technische Universität Darmstadt, Deutschland Universität Stuttgart, Deutschland Universität Trier, Deutschland Johann Wolfgang Goethe-Universität Frankfurt am Main, Deutschland Institut National de Recherche en Informatique et en Automatique (INRIA), Frankreich Centrum voor Wiskunde en Informatica (CWI), Niederlande Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Deutschland
Einbindung	Mitglied der Leibniz-Gemeinschaft Mitglied im Leibniz-Forschungsverbund Science 2.0
Verantwortlicher Redakteur	Prof. Raimund Seidel, Ph. D.
Redaktion	Michael Gerke, Dr. Shida Kunz, Dr. Michael Wagner
Mitarbeit	Dr. Marcel R. Ackermann, Heike Clemens, Dr. Michael Didas, Dr. Andreas Dolzmann, Jutka Gasiorowski, Michael Gerke, Dagmar Glaser, Oliver Hoffmann, Dr. Shida Kunz, Heike Meißner, Petra Meyer, Dr. Florian Reitz, Thomas Schillo, Raimund Seidel, Ph. D., Dr. Michael Wagner und die Autoren der Zusammenfassungen in Kapitel 6
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
ISSN	2199-1995
Online-Version	https://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	Juni 2020
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

Nach langen Vorarbeiten und Vorbereitungen war nun 2019 das erste Jahr, in dem die beliebte und viel verwendete Informatikliteraturdatenbank dblp gänzlich vom LZI betrieben wird. Ein sogenannter strategischer Sondertatbestand stellt auf Dauer beträchtliche Mittel zum Betrieb und Ausbau von dblp zur Verfügung. Der Personalausbau dafür an der neuen, offiziellen Außenstelle in Trier hat auch schon zu einer merklichen Verbesserung und Erweiterung des dblp-Service geführt. So werden z.B. offene Zitationsdaten verlinkt und auch offen zugängliche Versionen von Arbeiten. Es besteht natürlich weiterhin eine enge Kooperation mit der Universität Trier, an der dblp initiiert wurde.

Das Seminarprogramm deckte wieder ein weites Feld ab, von naturwissenschaftlichen Themen wie "Visual Computing in Materials Sciences" über traditionelle theoretische Informatikthemen wie "Computational Complexity of Discrete Problems" bis zu gesellschaftswissenschaftlichen Themen wie "Application-Oriented Computational Social Choice".

Auch unsere dritte Sparte, das Open-Access-Publishing, ist weiter gediehen mit nach wie vor steigenden Publikationszahlen und auch neuen Softwaresystemen für einen verbesserten Publikationsprozess.

2019 war ein gutes Jahr für Dagstuhl. Viele Details finden Sie auf den folgenden Seiten.

After long preparations, 2019 became the first year in which the popular and much used computer science publications database dblp was operated by the LZI alone. A so called strategic extraordinary item of expenditure provides a considerable amount of long-term funding for the operation and expansion of dblp. The accompanying increase in personnel at the new official branch office in Trier has already resulted in a marked improvement and expansion of the dblp service. For example, open citation data is linked, as are openly accessible versions of publications. Of course there is still a close cooperation with the University of Trier, where dblp was initiated.

The seminar program covered a wide range of topics again, from natural science topics like "Visual Computing in Materials Sciences" via classical theoretical computer science topics like "Computational Complexity of Discrete Problems" to social science topics like "Application-Oriented Computational Social Choice".

Our third branch—open access publishing—prospered as well, with still growing publication numbers and also with new software systems for an improved publication process.

2019 was a good year for Dagstuhl. You can find many details on the following pages.

Im Namen der Geschäftsführung

Prof. Raimund Seidel, Ph. D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

On behalf of the Managing Directors

Inhaltsverzeichnis
Contents

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 (siehe Fig. 1.1) jährlich mehr als 3 000 internationale Gäste.

Seit 2005 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 96 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 (see Fig. 1.1) hosts over 3,000 research guests each year from countries across the globe.

Since 2005, Schloss Dagstuhl has been a member of the Leibniz Association, a non-profit research consortium composed of 96 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 April 2020.
As of April 2020.

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 Gäste auch außerhalb des fachlichen Seminarprogramms kennenlernen und austauschen können.

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

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

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

Located in a 1760 build manor house in the idyllic countryside of northern Saarland at the heart of the tri-country region formed by Germany, France and Luxembourg, Schloss Dagstuhl offers visitors a unique working environment that encourages guests to interact with each other in tandem with daily life. Lounges, formal and informal dining areas, a world-class research library, and an impressive range of work and leisure rooms offer multiple possibilities for connecting one-on-one outside of the official conference rooms and meeting times.

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

Bibliographiedatenbank dblp

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

Mit mittlerweile mehr als 4,8 Millionen Publikationseinträgen ist dblp die weltweit größte offene Sammlung bibliographischer Daten in der Informatik. Der dblp-Dienst ist darauf ausgerichtet, Forscher bei ihrer täglichen Arbeit zu unterstützen, etwa bei der Literaturrecherche oder beim Bezug von elektronisch verfügbaren Volltexten. Dabei gilt dblp in der Informatik insbesondere als die Referenzdatenbank für qualitätsgesicherte, normierte Bibliographiedaten. Aber auch Forschungsförderer und Entscheidungsträger unterstützt dblp, etwa durch das Pflegen und öffentlich Verfügbarmachen von personalisierten Publikationsnachweisen. Durch den Betrieb von dblp leistet Schloss Dagstuhl einen weiteren Beitrag im Rahmen seiner Mission zur Förderung der Erkennung, Verbreitung und Umsetzung neuer Informatikerkenntnisse auf international anerkanntem Niveau.

Details über dblp finden sich in Kapitel 3.

Dagstuhl Publishing

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

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

Kapitel 4 stellt Dagstuhls Verlagswesen ausführlicher dar.

dblp computer science bibliography

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

Listing more than 4.8 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 bibliographies. 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 2019 / *News from 2019*

Das Team

Am Ende des Jahres 2019 beschäftigte Schloss Dagstuhl insgesamt 41 Vollzeitäquivalente bzw. 56 Mitarbeiter. Unsere Auszubildende hat im Herbst 2019 erfolgreich ihre Prüfung zur Hauswirtschafterin abgeschlossen. Drei Mitarbeiterinnen gingen in 2019/2020 in Mutterschutz bzw. Erziehungsurlaub. Dies erforderte bedingt Neueinstellungen, vorwiegend im Bereich Hauswirtschaft und Küche. Schloss Dagstuhl hat bereits 2018 begonnen, das

The Team

By the end of 2019, Schloss Dagstuhl had a total of 56 staff members corresponding to 41 full-time positions. Our trainee has successfully passed her examination as a housekeeper in the Fall of 2019. Three employees took maternity leave or parental leave in 2019/2020. This caused the hiring of temporary staff, especially in housekeeping and the kitchen. Already in 2018, Schloss Dagstuhl started to expand the dblp team as part of the planned expansion

Team von dblp im Rahmen der geplanten Erweiterung und Verstetigung von dblp als strategische Aufgabe von Schloss Dagstuhl aufzustocken. In 2019 kamen drei weitere Mitarbeiter in diesem Bereich hinzu.

Seminare und Workshops

In 2019 wurden 103 Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops gestellt. Dies liegt im oberen Bereich der langfristigen Tendenz zu etwa 100 Anträgen pro Jahr. Durch die rekordverdächtige Antragslage des Vorjahrs stieg die Anzahl der Seminare und Workshops in 2019 auf 73. Damit fanden 8 mehr als im Vorjahr und somit wieder über 70 Seminare und Workshops statt.

Von den mehr als 3 300 Gästen, die sich in Dagstuhl trafen, nahmen etwa 2 500 an Seminaren teil. Etwa 41% aller Seminarteilnehmer war zum ersten Mal in Dagstuhl, und fast ein Drittel der Teilnehmer an unserer Gastumfrage ordnete sich selbst als Junior-Wissenschaftler ein. Mehr als drei Viertel aller Seminarteilnehmer waren außerhalb von Deutschland beschäftigt.

Etwa 81% aller in 2019 stattgefundenen Seminare hatten mindestens eine Frau im Organisatorenteam, und rund 31 % aller Organisatoren waren Frauen. Der Frauenanteil unter allen Seminarteilnehmern betrug 24,1 %.

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

Rekordzahlen und neue Features für die Bibliographiedatenbank dblp

Seit nunmehr einem Jahr arbeitet das verstärkte dblp-Team am neu gegründeten LZI-Standort in Trier an der redaktionellen, technischen und wissenschaftlichen Verbesserung der Datenbank. Die neuen Möglichkeiten haben bereits zu sichtbaren Ergebnissen geführt. Mit mehr als 450 000 neu aufgenommenen Publikationen verzeichnete dblp das mit Abstand stärkste Wachstum in seiner 25-jährigen Geschichte. Diese Zahl entspricht mehr als 1 800 neuen Publikationen pro Arbeitstag, die von unseren Redakteuren geprüft und der passenden Bibliographie zugeordnet werden. Noch deutlicher fällt die Steigerung bei der Qualitätssicherung und Kuration bestehender Autoren-Bibliographien aus: Mit mehr als 55 000 Korrekturfällen hat sich die Anzahl gegenüber dem Vorjahr fast verdoppelt.

Doch auch die Nutzungszahlen des dblp-Dienstes zeigen einen deutlichen Anstieg. Inzwischen registrieren die dblp-Webseiten im Schnitt mehr als eine Million Seitenzugriffe pro Tag. Jeden Monat erhalten unsere Server Zugriffe von etwa 900 000 verschiedenen Nutzern aus aller Welt.

Gleichzeitig konnten die Leistungsmerkmale der Webdienste weiter ausgebaut werden. In 2019 feierten insbesondere offene Referenz- und Zitationsdaten ihr Debüt. Zudem zeichnet dblp inzwischen Open-Access-Publikationen sichtbar aus und stellt Links zu offenen Alternativen von Artikeln hinter der Bezahlschranke zur Verfügung. Um die Verwendung von dblp-Daten in der Forschung und deren Integration in externen Projekten zu erleich-

and perpetuation of dblp as a strategic task of Schloss Dagstuhl. 2019, three new staff members where added to the dblp team.

Seminars and Workshops

In 2019, 103 applications for Dagstuhl Seminars and Dagstuhl Perspectives Workshops were submitted. This is in line with the long-term trend of about 100 applications per year. Due to the very busy proposal submission period in the previous year, the number of seminars and workshops rose to 73 in 2019. With 8 more than the previous year, Schloss Dagstuhl was back to hosting more than 70 seminars and workshops this year.

Of the more than 3,300 guests who met in Dagstuhl, about 2,500 took part in seminars. About 41% of all seminar participants were in Dagstuhl for the first time, and almost a third of the participants in our guest survey classified themselves as junior scientists. More than three quarters of all seminar participants were employed outside Germany.

About 81% of all seminars held in 2019 had at least one woman in the team of organizers, and about 31 % of all organizers were women. Women accounted for 24.1 % of all seminar participants.

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

Record figures and new features for the dblp computer science bibliography

During the past year, the strengthened dblp team has been working on the editorial, technical and scientific improvement of the database at the newly founded LZI offices in Trier. The increased team size has already led to visible results. With more than 450,000 newly added publications in 2019, dblp had by far the strongest growth in its 25-year history. This figure corresponds to more than 1,800 publications per working day which have been checked and assigned to a bibliography by our editors. An even more significant increase has been seen for the quality assurance and curation of existing author bibliographies: With more than 55,000 edit cases the number has almost doubled compared to the previous year.

In addition, the usage statistics of dblp's web services also showed a clear increase in 2019. On average, the dblp website saw more than one million page impressions per day. Every month our servers received requests from about 900,000 different users from all over the world.

At the same time, new features have been added to dblp's web services. In particular, 2019 saw the debut of open reference and citation data in dblp. Furthermore, dblp now visibly marks open access publications and provides hyperlinks to open alternatives of articles behind the paywall. Since the end of 2019, the entire data set has been made available under the public domain CC0 licence in order to facilitate the use of dblp's data in research and its integration in external projects. This highlights our

tern, steht der gesamte Datensatz seit Ende 2019 unter der gemeinfreien CC0-Lizenz zur Verfügung. Dies soll unterstreichen, dass es sich bei dblp um ein allgemeines Gut der internationalen Informatik-Gemeinschaft handelt.

Mehr Details zu dblp finden sich in Kapitel 3.

Dagstuhl Publishing

Wie in den Vorjahren haben die Open-Access-Publikationsaktivitäten auch in 2019 starken Zuspruch bekommen. So wurden in den Konferenzbandreihe LIPIcs und OASIcs zusammen wieder über 1 300 Publikationen innerhalb eines Jahres veröffentlicht. Zudem gab es auch in 2019 wieder viele Bewerbungen von wissenschaftlichen Konferenzen zur Veröffentlichung des Konferenzbandes in der Serie LIPIcs. Zudem wurden erfolgreich ein neues Einreichungssystem und ein erweiterter Veröffentlichungsprozess eingeführt, welche die Arbeiten für Herausgeber, Editoren aber auch das Verlagsteam deutlich vereinfachen.

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

Dagstuhler Gespräche

Auch in 2019 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.

Für den 17. Mai 2019 konnte der Präsident der Gesellschaft für Informatik e.V., Prof. Dr. Hannes Federrath, als Vortragender gewonnen werden. Unter dem Titel „Was weiß das Internet über mich?“ erklärte er, welche Spuren man im Internet hinterlässt, und ging der Frage nach wie sich die Nutzer im Internet vor Profilbildung schützen können und welche Datenschutzrechte sie im Internet haben.

Noch eine weitere Veranstaltung der Reihe wurde durchgeführt: Am 24. November 2019 trug Prof. Dr. Katharina Zweig, Professorin für theoretische Informatik an der TU Kaiserslautern, zum Thema *Wie die Ethik in den Rechner kommt* vor. Sie ist Mitglied in verschiedenen Politikberatungsgremien, unter anderem im Koordinationsgremium des Netzwerks Verbraucherforschung und in der Enquete-Kommission Künstliche Intelligenz des Bundestages. Ihre Arbeit wurde mehrfach ausgezeichnet, unter anderem mit dem Communicator-Preis der DFG und des Stifterverbandes 2019. In ihrem Vortrag, der auch von ihrem neuen Buch *Ein Algorithmus hat kein Taktgefühl: Wo künstliche Intelligenz sich irrt, warum uns das betrifft und was wir dagegen tun können* inspiriert war, erklärte sie, an welchen Stellen beim Einsatz von Künstlicher Intelligenz die Antwort der Künstlichen Intelligenz durch menschliche Entscheidungen beeinflusst wird, und ermunterte die Anwesenden, sich bei solchen Entscheidungen einzumischen und sicherzustellen, dass diese ethisch getroffen werden.

understanding that dblp is a public good of the international computer science community.

More details 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 2019. So in the conference proceedings series LIPIcs and OASIcs together again over 1,300 publications were published within one year. Furthermore, LIPIcs again received and accepted proposals from several major scientific conferences. In addition, a new submission system and an enhanced publishing process have been successfully introduced, which significantly simplify the work for editors, publishers and the publishing team.

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

"Dagstuhler Gespräche"

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

The talk on May 17, 2019 was given by the former President of the Gesellschaft für Informatik e.V. (German Informatics Society) Prof. Dr. Hannes Federrath. Under the title "Was weiß das Internet über mich?" (What does the Internet know about me?), he explained what traces one leaves when on the Internet. He discussed how one can protect oneself from profiling and what data protection rights one has on the Internet.

One more event in this series took place: On November 24, Prof. Dr. Katharina Zweig, Professor for theoretical computer science at the TU Kaiserslautern, talked about the topic *Wie die Ethik in den Rechner kommt* (How ethics gets into the computer). She is a member of various political advisory bodies, inter alia the coordinating board of the network for consumer science and in the Bundestag's committee of inquiry on Artificial Intelligence. Her work has won several awards, for example the communicator-award of the DFG (German Research Foundation) and the Stifterverband (Donors' association for the promotion of humanities and sciences in Germany). In her talk, which was also inspired by her new book *Ein Algorithmus hat kein Taktgefühl: Wo künstliche Intelligenz sich irrt, warum uns das betrifft und was wir dagegen tun können* (An algorithm has no tact: Where artificial intelligence errs, why that matters to us, and what we can do about it), she explained where in the use of artificial intelligence

Beide Veranstaltungen waren rege besucht und lösten bei den Anwesenden eine rege Anteilnahme an den an den Vortrag anschließenden Gesprächen aus. Die Reihe wird im kommenden Jahr gewiss fortgesetzt werden.

Lehrerfortbildung

In Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisierte Schloss Dagstuhl 2019 zum 29. Mal eine Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet.

Zusammenarbeit mit dem Heidelberg Laureate Forum

Auch im Jahr 2019 gab es wieder eine Kooperation von Schloss Dagstuhl mit dem Heidelberg Laureate Forum[2] (HLF). Diese Veranstaltung bringt herausragende Mathematiker und Informatiker, nämlich Gewinner des ACM Turing Award, des Abelpreises, der Fields-Medaille, und des Nevanlinna-Preises, mit außergewöhnlich begabten jungen Wissenschaftlern aus aller Welt zusammen. Drei ausgewählte Teilnehmer des HLF 2019 erhielten in der Woche vor der siebten Ausgabe dieses Forums die Gelegenheit zur Teilnahme an dem Dagstuhl-Seminar „Application-Oriented Computational Social Choice“ (19381), ein weiterer konnte in der Woche danach am Dagstuhl-Seminar „Comparative Theory for Graph Polynomials“ (19401) teilnehmen.

Aufgrund des großen Erfolgs der Initiative erwarten wir eine Fortsetzung der Zusammenarbeit im Jahr 2020.

Spender und Förderer

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

DeepMind unterstützte das Dagstuhl Seminar „AI for the Social Good“ (19082) durch die Bereitstellung von 3 000 £[3] für Reisekosten von Teilnehmern aus Afrika.

2019 erhielt die Bibliothek von mehreren Verlagshäusern erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 672 Bände als Spende, darunter 631 Monographien des Springer-Verlags im Wert von 49 037 €.

Schloss Dagstuhl wurde 2019 durch verschiedene Kunstspenden unterstützt.

its answers are influenced by human decisions. She encouraged the audience to take a hand in such decisions and make sure they are made ethically.

Both talks were well attended and the discussions afterwards were lively. The Dagstuhler Gespräche will certainly see a continuation in the next year.

Teacher Training Program

In 2019, Schloss Dagstuhl hosted its teacher training course for the 29th time. This workshop is specifically designed for computer science and mathematics teachers teaching grades 11 and 12 in Saarland and Rhineland-Palatinate. It is organized in collaboration with the Landesinstitut für Pädagogik und Medien Saarland LPM (Saarland State Institute for Education and Media) and the Pädagogisches Landesinstitut Rheinland-Pfalz PL (Rhineland-Palatinate State Institute for Education).

Joint Outreach with the Heidelberg Laureate Forum

2019 saw another cooperation venture between Schloss Dagstuhl and the Heidelberg Laureate Forum[2] (HLF). The HLF brings winners of the ACM Turing Award, the Abel Prize, the Fields Medal, and the Nevanlinna Prize together with exceptionally talented young scientists from all over the world. Three participants were selected and participated in the Dagstuhl Seminar “Application-Oriented Computational Social Choice” (19381), taking place during the week before the seventh edition of the forum. Another participant was selected and participated in the Dagstuhl Seminar “Comparative Theory for Graph Polynomials” (19401) in the week after the forum.

Satisfied with the outstanding success of the initiative, we expect a continuation the cooperation in 2020.

Sponsors and Donors

Schloss Dagstuhl is grateful to its scientific guests and institutional colleagues for generous donations for the support of its center.

DeepMind supported the Dagstuhl Seminar “AI for the Social Good” (19082) with 3,000 £[3] for travel expenses of African participants.

The center’s research library received a large number of book donations from several publishing houses. The number of donated volumes totaled 672, including 631 monographs at the total value of 49,037 € donated by Springer Science+Business Media publishing house.

Schloss Dagstuhl was supported by various art donations in 2019.

[2] http://www.heidelberg-laureate-forum.org
[3] ∼3500€.

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 29 Forschern aus den USA eine Teilnahme an insgesamt 19 Seminaren ermöglicht werden. Insgesamt wurden dafür 49 293,25 $ Fördermittel ausgegeben. Die Förderung lief Ende September aus.

Baumaßnahmen und Renovierung

In 2019 wurde der Umbau des Personalraums fertiggestellt. Die Richtlinien des Arbeitsschutzes sind nun erfüllt und eine separate Nutzung von Umkleideraum und Personaltoiletten – nun getrennt für beide Geschlechter – wurde umgesetzt. Der Pausenraum wurde isoliert und mit einer kleinen Küchenzeile ausgestattet.

Um den Wärmeschutz im sogenannten Neubau im Sommer und im Winter zu verbessern, wurde 2019 damit begonnen, die Decken der Gästezimmer und der Verwaltung zu isolieren. Im Zuge dieser Maßnahme wurde als Vorbereitung eine Netzwerkverkabelung in der Decke angebracht, da die vorhandene Bausubstanz eine WLAN-Nutzung durch zentrale Zugriffspunkte sehr stark einschränkt.

In den Betriebsferien wurden u.a. vier Bäder saniert und notwendige Malerarbeiten im ganzen Institut durchgeführt.

Ende 2019 wurde eine motorgetriebene behindertengerechte Eingangstür eingebaut.

Ausstattung

Die Personalumkleide wurde mit neuen Doppelspinden ausgestattet, die Telefonanlage von Schloss Dagstuhl wurde teilweise erneuert. Vereinzelt wurden höhenverstellbare Schreibtische und neue Schreibtischstühle für unsere Mitarbeiter, insbesondere in Trier, angeschafft.

NSF Grant for Junior Researchers

Since 2013, Dagstuhl helps junior researchers based in the USA to participate in Dagstuhl seminars with funds provided through the grant opportunity[4] financed by the National Science Foundation (NSF).

In 2019, 29 US-based scientists were supported with a total amount of 49,293.25 $ and hence were able to participate in 19 Dagstuhl Seminars overall. The grant period ended in the end of September.

Construction Work and Renovation

The refurbishment of the staff rooms was completed in 2019. This improved occupational safety and facilitated a gender separated use of changing rooms and staff toilets. The break room was thermally insulated and fitted with a small kitchenette.

To improve the thermal insulation in the so called new building for summers and winters, thermal insulation of the ceilings of guest rooms and administrative offices was started in 2019. In the course of these works, network cabling was installed in the ceiling in a preparatory step, as the building fabric severely limits the usage of central WiFi access points.

During the vacation close-down, four bathrooms were refurbished and necessary paintwork all over the institute was performed.

Towards the end of 2019, a handicapped accessible automatic sliding front door was installed .

Facilities

The staff changing room was equipped with new double-lockers and the telephone system was partially upgraded.

A few height adjustable desks and office chairs were acquired for our staff, especially for the staff located in Trier.

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

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 26 Preisträger des Turing-Awards, der höchsten Auszeichnung, die im Bereich der Informatik auf internationaler Ebene verliehen wird.

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

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

Kapitel 6 enthält Zusammenfassungen der Dagstuhl-Seminare und Perspektiven-Workshops. Im Kapitel 14 sind alle Veranstaltungen, die 2019 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 26 winners of the ACM Turing Award, the highest achievable award within the international computer science community.

Dagstuhl's distinguished accomplishment is to have established pioneering, interdisciplinary seminars that have virtually become institutions themselves. Many of the topics addressed in-depth at Dagstuhl have subsequently developed into highly active research fields, resulting in some cases in DFG priority programs and other grant and funding programs. Dagstuhl Seminars often succeed in bringing together scientists from a range of research areas and disciplines whose work overlaps with respect to issues, methods and/or techniques, but who had never previously entered into constructive dialogue with one another. This especially applies to disciplines outside of the field of informatics. Key research areas for which in-depth collaboration with informatics specialists was initiated and consolidated at Dagstuhl include biology (since 1992) and sports (since 2006). The spectrum of seminar topics provides an excellent and broad overview of the areas currently under discussion in the informatics arena.

Each Dagstuhl Seminar is asked to contribute a record of the seminar proceedings in the form of a Dagstuhl Report. The report gives an overview of the seminar's program, talks, and results in a journal-like manner to allow for a high visibility and timely communication of its outcome. The periodical *Dagstuhl Reports* is published in one volume with 12 issues per year; each issue documents the Dagstuhl Seminars and Dagstuhl Perspectives Workshops of a given month. Dagstuhl Reports are openly accessible and can be downloaded at the Dagstuhl website.[5]

Chapter 6 contains a collection of the summaries of the 2019 Seminars and Perspectives Workshops. Chapter 14 provides a comprehensive list of all events that took place during the year under review, and a seminar program covering the coming 24 months is available on the Dagstuhl website.

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

Dagstuhl-Perspektiven-Workshops

2.2

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 2019 statt fanden, sind in Fig. 2.1 aufgelistet.

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

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

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

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

Dagstuhl Perspectives Workshops held in 2019 are listed in Fig. 2.1.

The results of the in-depth discussions of each workshop are presented in a manifesto detailing open issues and possible research perspectives in that specific field for the coming 5–10 years. Schloss Dagstuhl coordinates the targeted dissemination of this manifesto as research policy impulses to German and other European research donors and sponsors (EU, German Federal Ministry of Education and Research, DFG, etc.). Short versions of the manifestos are regularly presented in a forum of the *Informatik Spektrum* journal (published by Springer); full versions of the manifestos are published in our periodical Dagstuhl Manifestos[6].

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

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

The Role of Non-Monotonic Reasoning in Future Development of Artificial Intelligence http://www.dagstuhl.de/19072
Diversity, Fairness, and Data-Driven Personalization in (News) Recommender System http://www.dagstuhl.de/19482

Fig. 2.1
Dagstuhl Perspectives Workshops held in 2019.

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 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 2019 wieder ein hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 103 Anträgen in den Antragsrunden im Januar und Juni 2019. Der hohen Qualität der Anträge entsprechend, wurden etwa 56 % der eingereichten Anträge genehmigt. In den vergangenen 7 Jahren variierte die Rate der angenommen Anträge zwischen 56 % und 69 %(siehe Fig. 2.2).

Unter den 58 in 2019 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 6 Seminare bereits 2019 ausgerichtet werden, der Großteil wurde jedoch für das Seminar-Programm in 2020 eingeplant (hier und im Folgenden wird, sofern nicht anders angegeben, das Wort “Seminar” sowohl für Dagstuhl-Seminare als auch für Dagstuhl-Perspektiven-Workshops verwendet). Insgesamt 4 der 2019 genehmigten Seminare werden – aufgrund des Überhangs an großen Seminaren – in 2021 stattfinden.

Schloss Dagstuhl maintains the high quality of the Dagstuhl Seminar and Dagstuhl Perspectives Workshop series by identifying those proposals that promise a high potential to engage researchers – often from different disciplines – in scientific 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 2019, submitting 103 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2019 and June 2019 submission rounds. The quality of the proposals was excellent, resulting in a 56 % acceptance rate by Dagstuhl’s Scientific Directorate. In the previous seven years, proposal acceptance rates have tended to range between 56 % and 69 % (see Fig. 2.2).

Among the 58 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2019 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 2020 seminar program, although it was possible to schedule 6 of them already in 2019 (here and in the following, the word "seminar" is meant to include both Dagstuhl Seminars and Dagstuhl Perspektives Workshops, if not specified otherwise). A total of 4 seminars approved in 2019 will be held in 2021, as there are a lot of large seminars.

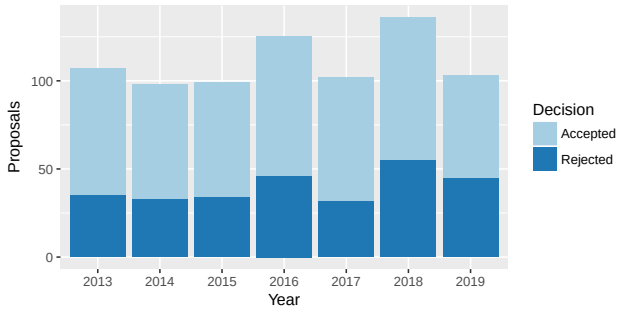

Fig. 2.2
Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2013–2019.

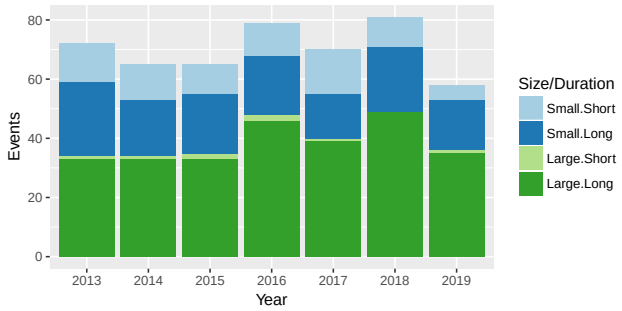

Fig. 2.3
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2013–2019.
Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

2.4 Seminar-Programm 2019 / *The Seminar Program in 2019*

In 46 von 48 Wochen, in denen das Tagungszentrum 2019 geöffnet war, fand mindestens ein Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop statt. In 27 Wochen waren es sogar zwei. In zwei 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. 2019 fanden insgesamt 73 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 46 of the 48 weeks the center was open in 2019. In 27 of those weeks, there were in fact two seminars in parallel. In the two remaining weeks, there were exclusively other events scheduled.

Since the guest house opened in 2012, it has been possible for the center to schedule two parallel seminars in any given week. Thus, there was an increase of seminars held since 2012 compared with the years before. Altogether, there were 73 Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2019. Fig. 2.4 shows the evolution in recent years.

2.5 Angaben zu Teilnehmern und Organisatoren / *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 2019 knapp die Hälfte (46 %, 1 081 von 2 369) der Wissenschaftler das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während weitere 16 % der Wissenschaftler an nur einem Seminar in den Jahren vorher teilgenommen hatten, weitere 9 % an zweien. Ein wenig andere Zahlen leiten sich aus unserer Gastumfrage ab. Hier ergibt sich, dass

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 2019, a bit less than half (1,081 of 2,369, or 46 %) of the researchers were first-time visitors to Dagstuhl. About an additional 16 % of the participating researchers had already attended one previous seminar in the years before, and another 9 % had already attended two. Slightly different numbers are obtained from our guest survey: About 41 % of

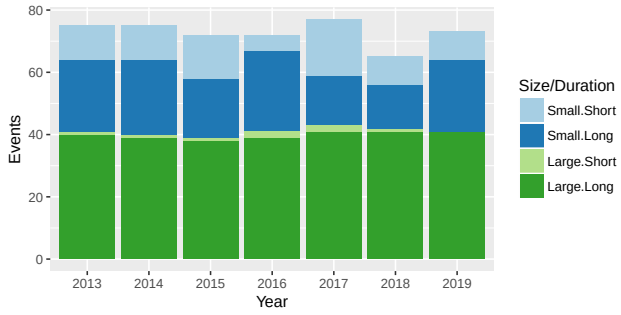

Fig. 2.4
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2013–2019.
Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

etwa 41 % der Antwortenden 2018 das erste Mal, 15 % 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 32 % der Gäste der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops in 2019, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein (siehe Fig. 2.5b). Diese ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern ist im Laufe der Jahre relativ konstant geblieben, was die Bemühungen des Zentrums zur Aufrechterhaltung der „Dagstuhl-Verbindung“ zwischen herausragenden jungen Wissenschaftlern und ihren erfahrenen Kollegen zeigt.

Mit 79 % war der Anteil von Seminarteilnehmern aus dem Ausland 2019 erneut sehr hoch. Das Diagramm in Fig. 2.5c zeigt die regionale Verteilung der Gäste für 2019 bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops. Mehr Details können Kapitel 13 entnommen werden.

In 2019 waren etwa 81 % aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt (siehe Fig. 2.6a). Der Anteil an weiblichen Seminarteilnehmern war mit 24,1 % höher als in den Jahren zuvor, was auch 2018 schon der Fall war (siehe Fig. 2.6b).

the responders were first-time visitors, an additional 15 % 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 32 % of 2019 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 and their senior colleagues.

At around 79 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2019. The chart in Fig. 2.5c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2019. For a detailed breakdown please refer to Chapter 13.

In 2019, 81 % 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 24.1 %, continuing the trend from 2018 (see Fig. 2.6b).

Themen und Forschungsgebiete 2.6 *Topics and Research Areas*

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

Im Folgenden sind beispielhaft einige thematische Schwerpunkte und dazugehörige Seminare aufgeführt. Die Aufzählung der Themen und Seminare hat keinen

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 2019. Neither the list of focal points nor the list of seminars is exhaustive. It

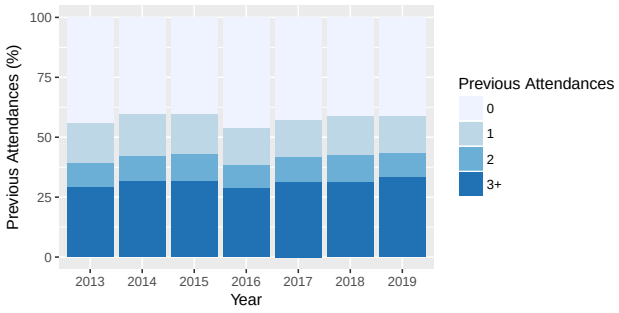

(a) Distribution of the number of previous attendances of participants, according to survey data.

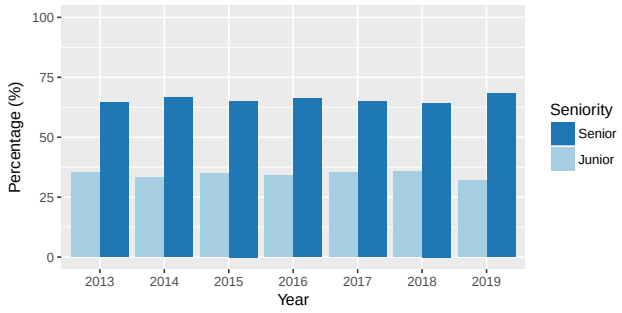

(b) Percentage of junior researchers, according to survey data.

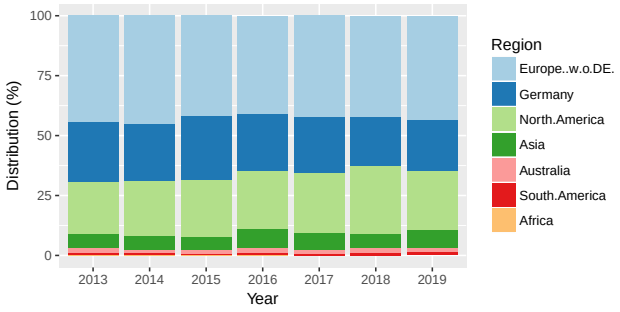

(c) Distribution of the origin of participants by region.

Fig. 2.5
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2013–2019.

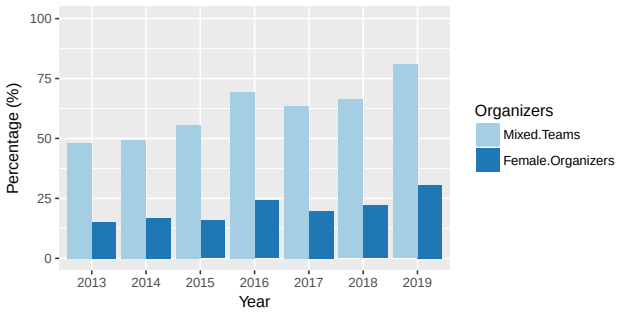

(a) Female organizers and mixed-gender organizer teams.

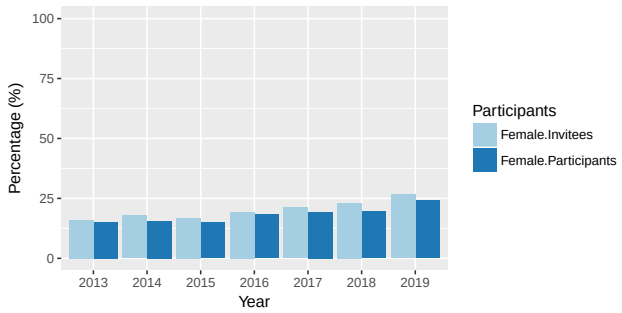

(b) Female invitees and participants.

Fig. 2.6
Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2013–2019.

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 2019.

Aktuell stark debattierte Themen waren natürlich deutlich vertreten, zum Beispiel Biotechnologie mit *Advances and Challenges in Protein-RNA Recognition, Regulation and Prediction* (19342), Sicherheit mit *Quantum Cryptanalysis* (19421) und *Biggest Failures in Security* (19451), oder autonomes Fahren mit *Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles?* (19132) und *Future Automotive HW/SW Platform Design* (19502).

Ein großer Schwerpunkt waren die Künstliche Intelligenz und Machine Learning und ihre Anwendungen und Konsequenzen, unter vielen Anderen mit *AI for the Social Good* (19082) und *Logic and Learning* (19361).

merely attempts to offer a brief insight into the multifarious scientific seminar program of 2019. Chapter 6, with the summary of the Seminars and Perspectives Workshops, provides a full overview of the 2019 scientific seminar program.

Topics that are heavily debated at the moment were of course represented considerably, for example Biotechnology with *Advances and Challenges in Protein-RNA Recognition, Regulation and Prediction* (19342), Security with *Quantum Cryptanalysis* (19421) and *Biggest Failures in Security* (19451), or autonomous driving with *Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles?* (19132) and *Future Automotive HW/SW Platform Design* (19502).

One focus was Artificial Intelligence and Machine Learning and their applications and consequences, with, among many others, topics like *AI for the Social Good* (19082) and *Logic and Learning* (19361).

Unter den Seminaren, die sich Themen aus dem Bereich der theoretischen Informatik gewidmet haben war ein Schwerpunkt die Komplexitätstheorie, etwa mit klassischen Untersuchungen wie *Algorithms and Complexity for Continuous Problems* (19341), aktuellen Trends wie *New Horizons in Parameterized Complexity* (19041), bis hin zu Anwendungen wie *Algorithms and Complexity in Phylogenetics* (19443).

Auch formale Methoden etwa mit *Bringing CP, SAT and SMT together: Next Challenges in Constraint Solving* (19062), auch in anderen Disziplinen wie *Verification and Synthesis of Human-Robot Interaction* (19081), waren gut vertreten.

Aber auch die Algorithmik mit *Beyond-Planar Graphs: Combinatorics, Models and Algorithms* (19092) und am anderen Ende des Spektrums das Software Engineering mit *BOTse: Bots in Software Engineering* (19471) kamen nicht zu kurz.

Ein weiterer Schwerpunkt war der Umgang mit Daten und Modellen bezüglich ihrer Visualisierung, von den Grundlagen wie *Visual Analytics of Multilayer Networks Across Disciplines* (19061), bis hin zu Anwendungen wie *Astrographics: Interactive Data-Driven Journeys through Space* (19262), aber auch Computergrafik im allgemeineren, etwa mit *3D Morphable Models* (19102).

Der Umgang mit großen Datenmengen war ebenfalls ein Thema, von *Theoretical Foundations of Storage Systems* (19111) bis Datenbanken bei *Data Series Management* (19282).

Nicht zuletzt wurden mit *Values in Computing* (19291) und *Ethics and Trust: Principles, Verification and Validation* (19171) auch Grundlegende Fragen der informatischen Ethik thematisiert.

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der in 2019 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.

Among the seminars which addressed topics from theoretical computer science, there was an emphasis on complexity theory, with classical topics like *Algorithms and Complexity for Continuous Problems* (19341), current trends like *New Horizons in Parameterized Complexity* (19041), up to applications like *Algorithms and Complexity in Phylogenetics* (19443).

Formal Methods were also well represented with topics like *Bringing CP, SAT and SMT together: Next Challenges in Constraint Solving* (19062). This extended to other disciplines like *Verification and Synthesis of Human-Robot Interaction* (19081).

But Algorithmics with topics like *Beyond-Planar Graphs: Combinatorics, Models and Algorithms* (19092) and – at the other end of the spectrum – Software Engineering with topics like *BOTse: Bots in Software Engineering* (19471) were not forgotten.

Another emphasis was on working with data and models and their Visualization, from the basics like *Visual Analytics of Multilayer Networks Across Disciplines* (19061), to applications like *Astrographics: Interactive Data-Driven Journeys through Space* (19262), but also computer graphics in gerenral, for example *3D Morphable Models* (19102).

Handling a lot of data was a topic as well, from *Theoretical Foundations of Storage Systems* (19111) to Databases with *Data Series Management* (19282).

Not least, *Values in Computing* (19291) and *Ethics and Trust: Principles, Verification and Validation* (19171) had a look at the ethical foundations of computer science.

This brief selection of seminars should not draw attention from the fact that each of the 2019 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 und Lehrerfortbildungen.

Further Event Types

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

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

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

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

Qualitätssicherung 2.8

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

Fig. 2.7 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2019 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 534 Fragebögen, welche die Meinung von etwa 61 % der 2 498 Teilnehmer repräsentieren. Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich.

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

Quality Assurance

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

Fig. 2.7 shows the satisfaction of responding participants in 2019 with regard to selected aspects of their stay. The results were compiled from 1,534 questionnaires, representing the responses of about 61 % of all 2,498 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 2019 konnte Schloss Dagstuhl die hohe Auslastung weitgehend halten. Es gab 2019 insgesamt 13 749 Gasttage, wobei 11 352 Gasttage auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Letztere Zahl ist der höchste Wert, der in den letzten fünf Jahren erreicht wurde. Insgesamt gab es bei den Gasttagen den Höchststand der letzten sechs Jahre. Es fanden im Berichtsjahr 112 Veranstaltungen mit insgesamt 3 309 Gästen statt. Weitere Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2019 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 2019 einschließlich Dagstuhl-Seminaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-

Utilization of the Center

Schloss Dagstuhl was able to uphold the high capacity utilization again in 2019. There were 13,749 overnight stays in total, with 11,352 overnight stays in Dagstuhl Seminars and Dagstuhl Perspectives Workshops. The latter number is higher than it was for the last five years. In total, the overnight stays in 2019 were at a peak compared to the last six years. The center hosted a total of 112 events with 3,309 guests in 2019. See Chapter 13 for further details.

Weekends were kept free in 2019, 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 2019, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and

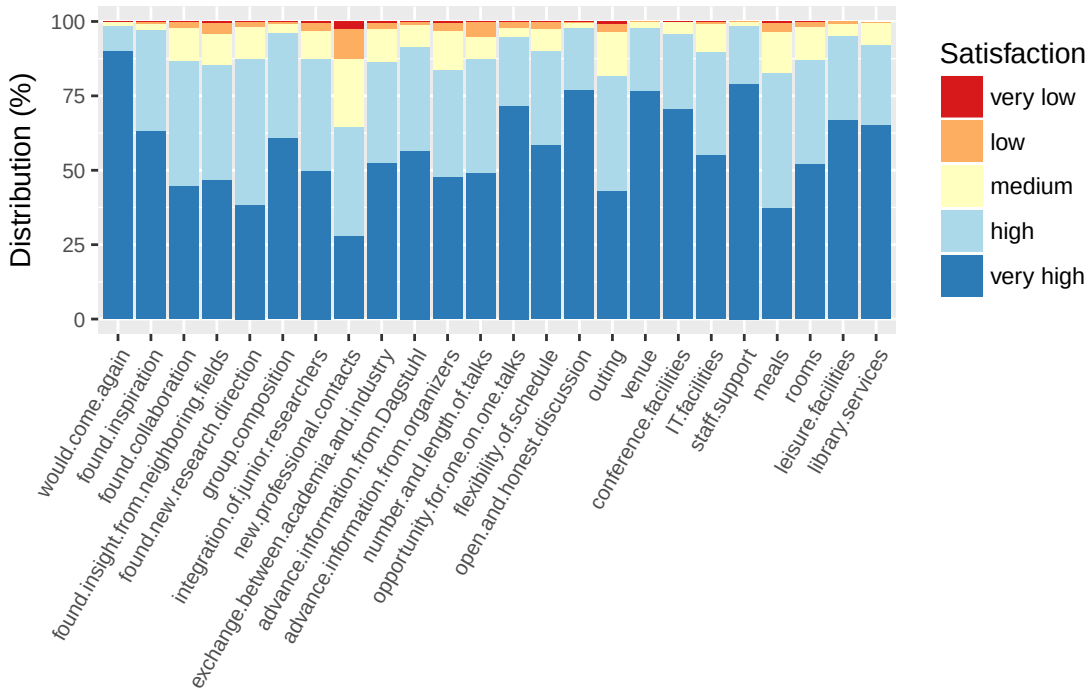

Fig. 2.7
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2019. According to survey results.

Seminaren und Veranstaltungen (z.B. Sommerschulen), bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Webseite ist ein Kalender[8] verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

host-only events such as meetings and summer schools can be found in Chapter 14. See the Schloss Dagstuhl website to view our calendar[8] of upcoming events and further information and materials on all events past, present and future.

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

3 Bibliographiedatenbank dblp
dblp computer science bibliography

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 25 Jahren einen substanziellen Beitrag durch die offene Bereitstellung qualitätsgeprüfter und aufbereiteter Publikationsdaten für die gesamte Informatik. Dabei unterstützt dblp die Informatik-Forschung auf gleich mehreren Ebenen, etwa durch:

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

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

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

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

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

[9] Google Scholar liefert zum Suchbegriff „dblp“ über 39 300 Treffer, Semantic Scholar liefert 13 300; im Einzelnen weisen SpringerLink ca. 3 800 Artikel, Elsevier ScienceDirect über 850 Artikel, die ACM Digital Library ca. 2 300 Artikel und IEEE Xplore über 2 500 Artikel nach.

[10] The search term “dblp” results in 39,300 hits at Google Scholar and 13,300 hits at Semantic Scholar; in particular, SpringerLink lists about 3,800 articles, Elsevier ScienceDirect lists more than 850 articles, the ACM Digital Library lists 2,300 articles, and IEEE Xplore lists more than 2,500 articles.

Schloss Dagstuhl und dblp

3.2 *Schloss Dagstuhl and dblp*

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

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

Dank der Finanzierung konnte das Team auf inzwischen 8 Vollzeitäquivalente aufgestockt werden, welche an der redaktionellen, technischen und wissenschaftlichen Verbesserung der Infrastruktur arbeiten. Dies hat bereits zu neuen Leistungsmerkmalen, einer spürbaren Steigerung der Zahl neu aufgenommener Publikationen und zu einer signifikanten Steigerung der Zahl der kuratierten Autorenbibliographien geführt (siehe Abschnitt 3.3–3.4).

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

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

Thanks to the increased funding the dblp team has grown to now 8 full-time equivalent staff members working on the editorial, technical, and scientific improvement of the infrastructure. The increased team size has already lead to new features, a noticeable improvement in the number of newly added publications and a significant increase of the number of curated author bibliographies (see Section 3.3–3.4).

[11] https://cirt.uni-trier.de/

dblp-Beirat | dblp Advisory Board
Prof. Dr. Hannah Bast University of Freiburg, Germany | *Chair*
Prof. Dr. Guillaume Cabanac Paul Sabatier University, Toulouse, France
Dr. Martin Fenner DataCite - International Data Citation Initiative e.V., Hannover, Germany
Prof. Dr. Silvio Peroni University of Bologna, Italy
Lydia Pintscher Wikimedia Deutschland - Association for the Promotion of Free Knowledge e.V., Berlin, Germany
Prof. Dr. Ruzica Piskac Yale University, New Haven, CT, USA
Prof. Dr. Rüdiger Reischuk University of Lübeck, Germany
Prof. Dr.-Ing. Ralf Schenkel University of Trier, Germany
Prof. Raimund Seidel, Ph.D. Saarland University, Saarbrücken, Germany

Fig. 3.1
The dblp Advisory Board in 2019.

Neue Funktionen und Verbesserungen

3.3

New features and improved services

Lizenzwechsel zu CC0

Seit Ende 2019 veröffentlicht dblp alle seine Daten unter der „CC0 1.0 Creative Commons Public Domain“ Lizenz. Dies betrifft insbesondere die täglichen und monatlichen Metadatenabzüge und die Web-APIs. Diese Änderung erleichtert die Nachnutzung unserer Daten in anderen offenen Datenprojekten erheblich: dblp-Daten können jetzt ohne Erlaubnis und für jeden Zweck (einschließlich kommerzieller Zwecke) verwendet werden. Eine verpflichtende Nennung von dblp als Quelle entfällt.

Unsere vorherige „ODC-BY 1.0 Open Data Commons Attribution“ Lizenz wurde 2011 gewählt, als Creative-Commons-Lizenzen noch nicht die beste Wahl für Datenveröffentlichungen waren. Seitdem sind sie jedoch durch die CC 4.0-Lizenzen überholt. Darüber hinaus erschwerte die zwingende Forderung einer Quellennennung die Nachnutzung und Integration, insbesondere in Linked Open Data-Szenarien wie WikiData.[12] Um die Kompatibilität mit bestehenden Integrationen zu gewährleisten werden wir auf absehbare Zeit weiterhin ODC-BY als Sekundärlizenz anbieten.

Offene Referenz- und Zitationsdaten

Eine der am häufigsten angefragten Funktionserweiterungen von dblp war in den letzten Jahren ohne jeden Zweifel das Hinzufügung von Zitationsdaten. Obwohl eine solche Funktion mit der Ressourcen von dblp bisher nicht möglich war, hat sich dies im Laufe des Jahres 2019 geändert: Ein großer Teil der vorhandenen Zitationsdaten wurde der Öffentlichkeit zur Wiederverwendung zugänglich gemacht. Dies gelang hauptsächlich dank der Bemühungen von Crossref[13], OpenCitations[14] und der „Initiative for Open Citations“ (I4OC).[15]

Unter Verwendung dieser offen Datenquellen konnten wir eine neue „Referenzen & Zitationen“-Ansicht für jede Publikation in dblp erstellen, die einen DOI besitzt.[16] Aktuelle Statistiken (Stand: November 2019) zeigen, dass zumindest eine teilweise Referenzliste sich für 51,6% aller Publikationen im dblp finden lässt, und dass 45,7% aller Publikationen im dblp zumindest eine zitierende Arbeit auflisten. Dies ist zwar schon recht beeindruckend, dennoch befindet sich die neue Detailansicht in vielerlei Hinsicht noch in der Entwicklung. Zudem gilt zu bedenken, dass eine Reihe wichtiger Informatik-Verlage (wie etwa IEEE und Elsevier) noch immer keine offenen Zitationsdaten bereitstellen. Daher besteht eine systematische Verzerrung bei der Verfügbarkeit dieser Daten.

License change to CC0

Since late 2019, all of dblp's data is released as a public commodity under the “CC0 1.0 Creative Commons Public Domain” license. This affects all metadata releases, in particular the daily and monthly data dumps and data retrieved from the web APIs. This change made it much easier to reuse our data in other open data projects. In a nutshell, dblp data can now be used without asking permission, for any purpose (including commercial purposes), and even without attributing the data to dblp.

Our previous “ODC-BY 1.0 Open Data Commons Attribution” license was selected as a fitting license back in 2011, when Creative Commons licenses were not yet a best choice for handling data publications. But since then, it has become somewhat superseded by CC 4.0 licenses. Furthermore, and probably most important, the attribution requirement made reuse and data integration more difficult in linked open data scenarios like WikiData.[12] To ensure compatibility with all existing integrations of dblp data, we will continue to provide ODC-BY as a secondary license for the foreseeable future.

Open references and citations

One of the most frequently requested features in the past years has been the addition of citations and references to dblp. While such a feature used to be infeasible within the resources of dblp, this has changed during the course of 2019: A huge chunk of the existing bibliographic citation data has been opened up to the public for reuse, mainly thanks to the efforts of Crossref[13], OpenCitations[14], and the “Initiative for Open Citations” (I4OC).[15]

Using these openly available reference and citation data sources, we were able to build a new “references & citations” details view for each publication in dblp that is assigned with a DOI.[17] Current statistics (as of November 2019) show that you can find at least a partial reference list for 51.6% of all publications in dblp, and that 45.7% of all publications in dblp list at least one citing paper. While this is already quite impressive, in many ways, the reference and citation details are still a work in progress. Also, one should keep in mind that a number of important publishers in computer science (such as IEEE and Elsevier) are still not supporting open citation data. Hence, there is a systemic bias in the availability of such data.

[12] https://wikidata.org
[13] https://www.crossref.org
[14] https://opencitations.net
[15] https://i4oc.org
[16] Bitte beachten Sie, dass die aktuellen offenen Datenquellen keine Informationen über Publikationen ohne DOI besitzen.
[17] Please note that using the open data sources, it is currently not possible to retrieve citation data for publications without a DOI.

Open-Access-Versionen von Artikeln

Die meisten der fünf Millionen im dblp-Datensatz enthaltenen Hyperlinks verweisen auf Artikelseiten innerhalb der digitalen Bibliothek eines Verlags. Eine wachsende Zahl von Verlagen hat das Open-Access-Publikationsmodell adaptiert und ermöglicht damit die kostenlose und barrierefreie Verbreitung von Forschungsergebnissen. Im Jahr 2019 haben wir damit begonnen, die Hyperlinks solcher Artikel mit einer speziellen, orangefarbenen Plakette zu kennzeichnen, um deren Verfügbarkeit zu signalisieren.[18]

Die meisten Informatik-Verlage verlangen jedoch immer noch ein Abonnement oder eine Gebühr für den Zugang zu Forschungsartikeln. Aufgrund der Art und Weise, wie dblp seine Metadaten sammelt, sind diese kostenpflichtigen Seiten in der Regel die einzigen, die in dblp verlinkt sind. Dank der neuen Integration von Unpaywall[19]-Daten können wir nun in vielen Fällen eine Alternative abseits der Bezahlschranke anbieten. Unpaywall ist eine Sammlung von Hyperlinks zu frei zugänglichen Versionen wissenschaftlicher Publikationen. Dazu gehören unabhängig archivierte oder erweiterte Autorenkopien (z.B. auf dem Webserver eines Instituts) sowie Preprints in offenen Repositorien. Gegenwärtig listet Unpaywall etwa 25 Millionen frei zugängliche Dokumente aus allen wissenschaftlichen Disziplinen, wobei eine beträchtliche Anzahl davon Informatik-Publikationen sind. Die Links zu diesen freien Versionen werden auf der dblp-Website neben der üblichen Liste externer Links zusammen mit einer kleinen, grünen Plakette angezeigt.

Verknüpfung mit externen Kennungen

Bis Ende 2019 waren bereits mehr als 77 000 dblp-Autorenbibliographien mit insgesamt mehr als 150 000 redaktionell geprüften, externen Kennungen verknüpft. Dies ist eine Zunahme von 55,2% gegenüber dem Vorjahr. Zu diesen Kennung gehören zentrale Normkennungen (z.B. ORCID, WikiData oder ISNI), Kennungen innerhalb von Nationalbibliotheken (z.B. LOC und GND) und digitalen Verlagsbibliotheken (wie ACM und IEEE), sowie weit verbreitete, kommerzielle und proprietäre Profile wie Google Scholar, LinkedIn oder Twitter.

Die Verknüpfung mit externen Kennungen ist aus einer Reihe von Gründen sehr nützlich. Zum einen helfen diese bei der Kuratierung von Autorenbibliographien, bei denen Metadaten wie Autorennamen mehrdeutig sind. Widersprüchliche Kennungen decken homonyme und synonyme Fälle auf, die sonst unbemerkt blieben. Des weiteren ermöglichen externe Ressourcen unseren Nutzern, ein vollständigeres Bild über die gelisteten Autoren zu erhalten. Dies gilt insbesondere für Autoren, die an inter- oder transdisziplinären Themen arbeiten. Verknüpfte Ressourcen dienen zudem als Datenquellen zur Erweiterung und Verbesserung der in dblp verfügbaren Informationen und tragen zum Aufbau eines semantisch ausdrucksstarken Linked Open Data-Netzwerkes bei.

Open access article versions

Most of the five million hyperlinks contained in the dblp data set point to article landing pages within a publisher's digital library. A growing number of publishers have adopted the open access model of publishing, thereby allowing the dissemination of research results free of cost and without any access barrier. In 2019, we have begun to mark such hyperlinks on the dblp website with a special orange badge signaling their availability.[20]

However, most publishers in computer science do still demand an active subscription or a fee for access to research articles. Due to the way dblp collects its metadata, these paywalled document locations are usually the only ones that are listed in dblp. But thanks to the new integration of Unpaywall[19] data, in many cases, we are now able to provide an open alternative. Unpaywall is a collection of hyperlinks to openly accessible versions of scholarly publications. This includes independently archived author copies (e.g., hosted on an institute's web server), preprints in open repositories, or open full versions of paywalled extended abstracts. Currently, Unpaywall lists about 25 million openly accessible documents across all scholarly disciplines, with a fair number of them being computer science publications. The links to those unpaywalled versions are displayed on the dblp website among the usual list of external links together with a small green badge.

Linking with external identifiers

By end of 2019, there are more than 77,000 dblp author bibliographies linked with more than 150,000 editorially checked external identifiers. This is an increase of 55.2% compared to the previous year. These identifiers include central authority control schemes (like ORCID, WikiData, or ISNI), identifiers within national libraries (e.g., LOC and GND) and publisher's digital libraries (such as ACM and IEEE), as well as prominent commercial and proprietary profiles like Google Scholar, LinkedIn, or Twitter.

There are a number of reasons why linking with external identifiers is very useful. First, due to their uniqueness, external IDs help when curating author bibliographies where the usual metadata like author names might still be ambiguous. Conflicting identifiers are particularly helpful to uncover homonym and synonym cases which might otherwise go unnoticed. Second, external resources allow our users to get more complete information about an author. This especially applies in cases where an author is working on inter- or cross-disciplinary topics which are not well covered by dblp alone. Linked resources serve as data sources for expanding and improving the information available in dblp and contribute to building a semantically meaningful linked open data network.

[18] Bitte beachten Sie, dass diese Auszeichnung noch in Arbeit ist und dass es noch viele frei zugängliche Artikel in dblp gibt, die derzeit unerkannt sind.
[19] https://unpaywall.org
[20] Please note that this badge is still work in progress, and that there are still plenty of openly accessible articles in dblp that go unrecognized.

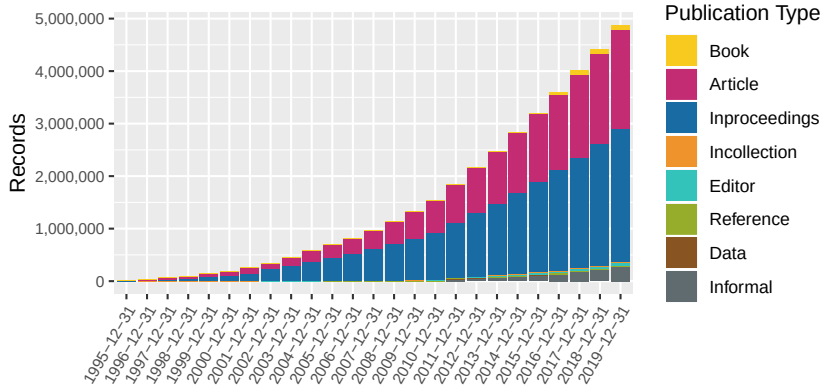

(a) Total number of records by year and type

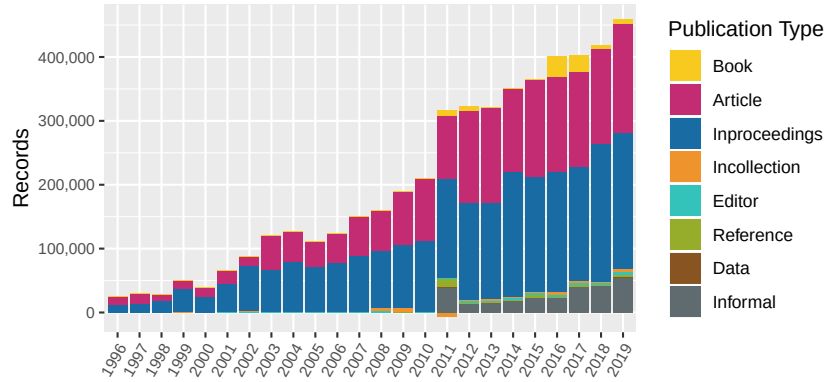

(b) New records by year and type

Fig. 3.2
Development of the dblp data stock.

Statistiken der Datenakquise 3.4 *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.

Dank der Erweiterung des dblp-Teams wurde im Laufe des Jahres 2019 eine Rekordanzahl von über 450 000 neuen Publikationseinträgen aufgenommen; dies entspricht mehr als 1 800 neuen Publikationen pro Arbeitstag. Ende Dezember 2019 indexierte dblp damit bereits mehr als 4,8 Millionen Publikationen aus den verschiedenen Teilgebieten der Informatik. Die neu aufgenommenen Einträge verteilen sich zu 46,1% auf Konferenzbeiträge, zu 37,1% auf Journalartikel, zu 11,9% auf Preprints und „graue“ Literatur, sowie zu 4,9% auf andere Publikationstypen wie etwa Monographien und Dissertationen.

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.

Thanks to the increased size of the dblp team, the dblp database grew by more than 450,000 publication records. This is the largest figure ever achieved in the history of dblp and corresponds to more than 1,800 new records for each working day of the year. By the end of December 2019, more than 4.8 million publications have been indexed by dblp. This year's new records consist of 46.1% conference papers, 37.1% journal articles, 11.9% preprints and "grey" literature, and 4.9% further publication types like monographs and PhD theses.

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

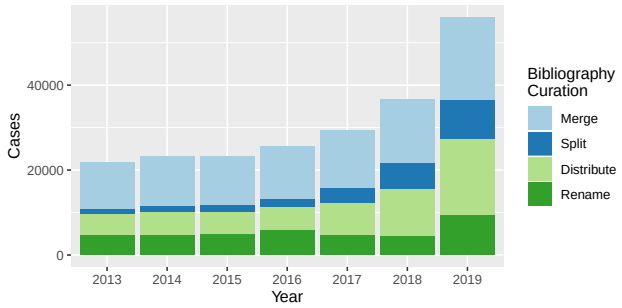

Fig. 3.3

Curation of existing dblp author bibliographies. The figures give the number of distinct edit cases (measured between the first and the last day of every given year) where a dblp team member manually corrected the assignment of publications within dblp author bibliographies. We distinguish between four curation cases: *Merge* = Two or more synonymous bibliographies have been merged into a single bibliography. *Split* = A single, homonymous bibliography has been split into two or more bibliographies. *Distribute* = A mixed case where records from two or more bibliographies have been redistributed between two or more bibliographies. *Rename* = A case where no actual publications have been reassigned, but the surface form of the author name(s) of a bibliography have been corrected or improved.

Statistiken der Datenkuration 3.5 *Data Curation Statistics*

Ein Hauptziel unserer intensiven Datenpflege ist es sicherzustellen, dass unsere Autorenbibliographien so korrekt und vollständig wie möglich sind. Das bedeutet, dass alle Publikationen eines Autors in nur einer einzigen Bibliographie aufgeführt sein sollen und dass diese Bibliographie auch nur Publikationen des spezifischen Autors listen soll. Es kann ziemlich schwierig sein, dies zu gewährleisten, und trotz unserer Bemühungen ordnen wir regelmäßig Publikationen einer falschen Bibliographie zu. Aus diesem Grund überprüft unser Redaktionsteam ständig unsere Daten und korrigiert solche Fehler.

Während spezielle Algorithmen dem Team helfen, solche Defekte aufzudecken, werden Korrekturen immer auf der Grundlage der Entscheidung eines menschlichen Kurators durchgeführt. Dies ist notwendig, da die verfügbaren Metadaten in der Regel nicht genügend Informationen enthalten, um eine hochpräzise automatisierte Lösung zu erlauben. Oft ist eine manuelle Recherche unter Berücksichtigung externer Ressourcen erforderlich.

Dank der Vergrößerung des dblp-Teams konnten wir noch mehr Arbeit in die Identifizierung und Korrektur fehlerhafter Bibliographien investieren. Daher ist die Gesamtzahl der gelösten Fehlerfälle in 2019 gegenüber dem Vorjahr um 52,3% auf über 55 000 Fälle gestiegen. Fig. 3.3 zeigt die Anzahl der gelösten Fehlerfälle in den letzten Jahren.

One main goal of the intensive data curation at dblp is to ensure that our author bibliographies are as correct and complete as possible. This means that all publications of a person should be listed in a single bibliography, and that a bibliography should only list publications from that specific author. It can be quite difficult to ensure this, and despite our best efforts, we regularly assign publications to the wrong bibliography. Because of this, our editorial team constantly checks our data and corrects such defect.

While specialized algorithms help our team to uncover and identify the nature of defect in our data, corrections are always executed based on the decision made by a human curator. This is necessary since the available metadata usually does not carry enough information to allow for a highly precise automated solution, and often requires a manual investigation taking external resources into account.

Due to the increased size of the dblp team, in 2019, we have been able to spent even more effort into identifying and correcting defective author bibliographies and publication misassignment. Hence, the total number of resolved defect cases has increased by 52.3% to over 55,000 cases per year. Figure 3.3 shows the number of resolved defect cases during the past few years.

	Trier 1		Trier 2		Dagstuhl		Total		
	2018	2019	2018	2019	2018	2019	2018	2019	%
user sessions (visits) per day	31 530	31 024	3 233	1 808	11 483	24 994	46 247	57 827	+25.0
page views per day	618 067	735 190	20 208	22 761	202 301	326 053	840 577	1 084 005	+29.0
page views per user session	19,6	23,7	6,2	12,6	17,6	13,0	18,2	18,7	+3.1
distinct users (IPs) per month	451 769	466 015	27 448	12 963	197 270	424 106	676 489	903 085	+33.5
data served per month	1 535,0 GB	2 114,1 GB	72,6 GB	89,6 GB	469,7 GB	821,3 GB	2 077,3 GB	3 025,0 GB	+45.6

Fig. 3.4
Average usage of the three dblp web servers. Trier 1 = dblp.uni-trier.de, Trier 2 = dblp2.uni-trier.de, Dagstuhl = dblp.dagstuhl.de

Nutzungsstatistiken 3.6 *Usage Statistics*

Im Jahr 2019 wurden vom dblp-Team drei offizielle dblp-Server geführt. Die Daten dieser Server werden täglich aktualisiert und miteinander synchronisiert:

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

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

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

Insgesamt konnten die Nutzungszahlen in 2019 gegenüber dem Vorjahr deutlich gesteigert werden. So bedienen die dblp-Webserver inzwischen täglich mehr als eine Million Seitenaufrufe. Fig. 3.4 fasst die durchschnittliche Nutzung aller drei dblp-Server zusammen. Diese Statistiken ignorieren die Zugriffe, die durch bekannte Bot- und Crawler-Software verursacht wurden.

In 2019, three official dblp web servers were updated and synchronized on a daily basis:

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

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

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

Overall, the total usage figures in 2019 significantly improved when compared to the previous year. In particular, the dblp web servers now serve more than one million page impressions per day. Figure 3.4 shows the average usage of all three servers. These figures ignore the traffic caused by known bots and crawlers.

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 um die wissenschaftliche Zeitschrift *LITES* und die Serie *DARTS*, in der Forschungsartefakte veröffentlicht werden.

Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports*[21] dokumentiert, um eine Zitation der Seminare im wissenschaftlichen Kontext zu ermöglichen. Zudem bietet sie auch den Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift erscheint seit 2011 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 2019 wurde für 75 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 Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*[22] ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop basiert. Für jedes Buch ist ein Antrag notwendig, der vom wissenschaftlichen Direktorium (welches als Herausgebergremium verantwortlich ist) begutachtet und freigegeben werden muss. In 2019 wurde kein Buch in der Reihe veröffentlicht.

[21] https://www.dagstuhl.de/dagrep
[22] https://www.dagstuhl.de/dfu

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* and by the DARTS series which aims at publishing research artifacts.

Dagstuhl Reports

All Dagstuhl Seminars and Dagstuhl Perspectives Workshops are documented in the periodical *Dagstuhl Reports*[21] which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

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

In 2019, 75 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 Follow-Ups

The *Dagstuhl Follow-Ups*[22] book series is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl Perspectives Workshop. Each book requires a proposal, which is reviewed and finally approved by the Scientific Directorate (which is in charge as editorial board). In 2019, no volume was published in the series.

Dagstuhl Publishing

Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*[23] 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 2019 wurde eine Ausgabe mit fünf Manifestos veröffentlicht (siehe Fig. 4.1).

DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*[24] werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2019 wurde die fünfte Ausgabe mit zwei Heften und insgesamt 20 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.

Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*[23] in an Open Access manner. The Scientific Directorate (see Fig. 11.4) acts as the editorial board of the journal. In 2019 one volume with five manifestos was published (see Fig. 4.1).

DARTS: Dagstuhl Artifacts Series

The *DARTS* series[24] publishes evaluated research data and artifacts. It is organized as a periodical. In 2019, one volume containing two issues with 20 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 focuses on conference proceedings publications.

[23] https://www.dagstuhl.de/dagman
[24] https://www.dagstuhl.de/darts
[25] https://www.dagstuhl.de/15452
[26] https://www.dagstuhl.de/16041
[27] https://www.dagstuhl.de/16111

Research Directions for Principles of Data Management (Dagstuhl Perspectives Workshop 16151)
Dagstuhl Manifestos, Volume 7, Issue 1, pp. 1-29, https://doi.org/10.4230/DagMan.7.1.1
based on Dagstuhl Perspectives Workshop 16151 https://www.dagstuhl.de/16151

QoE Vadis? (Dagstuhl Perspectives Workshop 16472)
Dagstuhl Manifestos, Volume 7, Issue 1, pp. 30-51, https://doi.org/10.4230/DagMan.7.1.30
based on Dagstuhl Perspectives Workshop 16472 https://www.dagstuhl.de/16472

Tensor Computing for Internet of Things (Dagstuhl Perspectives Workshop 16152)
Dagstuhl Manifestos, Volume 7, Issue 1, pp. 52-68, https://doi.org/10.4230/DagMan.7.1.52
based on Dagstuhl Perspectives Workshop 16152 https://www.dagstuhl.de/16152

Present and Future of Formal Argumentation (Dagstuhl Perspectives Workshop 15362)
Dagstuhl Manifestos, Volume 7, Issue 1, pp. 69-95, https://doi.org/10.4230/DagMan.7.1.69
based on Dagstuhl Perspectives Workshop 15362 https://www.dagstuhl.de/v

From Evaluating to Forecasting Performance: How to Turn Information Retrieval, Natural Language Processing and Recommender Systems into Predictive Sciences (Dagstuhl Perspectives Workshop 17442)
Dagstuhl Manifestos, Volume 7, Issue 1, pp. 96-139, https://doi.org/10.4230/DagMan.7.1.96
based on Dagstuhl Perspectives Workshop 17442 https://www.dagstuhl.de/17442

Fig. 4.1
Manifestos published in the 2019 volume of the journal *Dagstuhl Manifestos*.

Dagstuhl Publishing

OASIcs: OpenAccess Series in Informatics

Die *OASIcs*-Reihe[28] 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 2019 wurden 9 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.3.

OASIcs: OpenAccess Series in Informatics

The *OASIcs* series[28] 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 2019, Dagstuhl published 9 *OASIcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 4.3.

[28] https://www.dagstuhl.de/oasics

Prof. Dr. Daniel Cremers TU Munich, Germany
Prof. Dr. Barbara Hammer Bielefeld University, Germany
Prof. Dr. Marc Langheinrich University of Lugano, Switzerland
Prof. Dr. Dorothea Wagner Karlsruhe Institute of Technology, Germany | Chair

Fig. 4.2
OASIcs Editorial Board.

Vol. 66 | 2018 Imperial College Computing Student Workshop (ICCSW 2018) https://www.dagstuhl.de/dagpub/978-3-95977-097-2
Vol. 67 | 9th Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU 2018) https://www.dagstuhl.de/dagpub/978-3-95977-091-0
Vol. 68 | Workshop on Autonomous Systems Design (ASD 2019) https://www.dagstuhl.de/dagpub/978-3-95977-102-3
Vol. 69 | 2nd Symposium on Simplicity in Algorithms (SOSA 2019) https://www.dagstuhl.de/dagpub/978-3-95977-099-6
Vol. 70 | 2nd Conference on Language, Data and Knowledge (LDK 2019) https://www.dagstuhl.de/dagpub/978-3-95977-105-4
Vol. 72 | 19th International Workshop on Worst-Case Execution Time Analysis (WCET 2019) https://www.dagstuhl.de/dagpub/978-3-95977-118-4
Vol. 73 | 4th International Workshop on Security and Dependability of Critical Embedded Real-Time Systems (CERTS 2019) https://www.dagstuhl.de/dagpub/978-3-95977-119-1
Vol. 74 | 8th Symposium on Languages, Applications and Technologies (SLATE 2019) https://www.dagstuhl.de/dagpub/978-3-95977-114-6
Vol. 75 | 19th Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2019) https://www.dagstuhl.de/dagpub/978-3-95977-128-3

Fig. 4.3
OASIcs volumes published in 2019.

Dagstuhl Publishing

LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe[29] 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 Susanne Albers, Michael Mitzenmacher, Madhavan Mukund und Reinhard Wilhelm sind 2019 ausgelaufen. Alle haben als langjährige Mitglieder des Herausgebergremiums eine wichtige Rolle in der Entwicklung der Serie gespielt. Für diese Verdienste möchten wir uns an dieser Stelle herzlich bedanken.

Luke Ong, Meena Mahajan, Mikolaj Bojanczyk, Roberto Di Cosmo und Dieter van Melkebeek wurden in einem anonymen Wahlverfahren innerhalb des Herausgebergremiums neu in das Gremium gewählt. Zudem wurde Luca Aceto als Vorsitzender des Herausgebergremiums für zwei weitere Jahre bestätigt. Siehe auch Fig. 4.4.

In 2019 wurden Tagungsbände von 29 Konferenzen mit insgesamt 1208 Artikeln veröffentlicht; siehe Fig. 4.5 und 4.6.

Auch im zurückliegenden Jahr 2019 gab es wieder viele Anträge bei LIPIcs, womit die große Nachfrage aus den Vorjahren fortgesetzt wurde. In Fig. 4.7 sind alle Konferenzen aufgelistet, deren Anträge 2019 bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischweise 5 Jahre) eingegangen wurde. Vier dieser Konferenzen haben erstmals einen Antrag bei LIPIcs gestellt. Die anderen Konferenzen haben bereits vorher mit LIPIcs kooperiert.

LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series[29] 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 Susanne Albers, Michael Mitzenmacher, Madhavan Mukund, and Reinhard Wilhelm ended in 2019. All served as members of the editorial board for several years and played an important role for the development of the series. We would like to take this opportunity to thank them for their extraordinary dedication.

Luke Ong, Meena Mahajan, Mikolaj Bojanczyk, Roberto Di Cosmo, and Dieter van Melkebeek were voted in anonymous voting within the editorial board as new members of the editorial board. In addition, Luca Aceto was confirmed as chair of the editorial board for a further two years. See also Fig. 4.4.

The series published the proceedings of 29 major conferences with more than 1208 articles in total in 2019; see Fig. 4.5 and 4.6.

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2019 has again seen several applications for LIPIcs, continuing the high interest from the previous years. Fig. 4.7 lists all conferences that have been accepted in 2019 for a cooperation covering several years (typically 5 years). Four of these conferences have submitted a proposal to LIPIcs for the first time. The other conferences have already cooperated with LIPIcs in the past.

[29] https://www.dagstuhl.de/lipics

Prof. Dr. Luca Aceto Gran Sasso Science Institute, Italy and Reykjavik University, Iceland | Chair	Prof. Michael Mitzenmacher, Ph. D Harvard University, US | tenure ended in May 2019
Prof. Dr. Susanne Albers Technical University Munich, Germany | tenure ended in May 2019	Prof. Madhavan Mukund, Ph. D. Chennai Mathematical Institute, India | tenure ended in May 2019
Prof. Dr. Christel Baier Technische Universität Dresden, Germany	Prof. Dr. Anca Muscholl LaBRI and University Bordeaux, France
Prof. Dr. Mikolaj Bojanczyk University of Warsaw, Poland | tenure started in June 2019	Prof. Dr. Luke Ong University of Oxford, United Kingdom | tenure started in June 2019
Prof. Dr. Roberto Di Cosmo INRIA and University Paris Diderot, France | tenure started in June 2019	Dr. Catuscia Palamidessi INRIA, France
Prof. Dr. Javier Esparza Technical University Munich, Germany	Prof. Dr. Thomas Schwentick TU Dortmund, Germany
Prof. Dr. Meena Mahajan Institute of Mathematical Sciences, India | tenure started in June 2019	Prof. Raimund Seidel, Ph. D. Saarland University, Germany
Prof. Dieter van Melkebeek, Ph. D. University of Wisconsin-Madison, USA | tenure started in June 2019	Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Saarland University, Germany | tenure ended in May 2019

Fig. 4.4
LIPIcs Editorial Board.

Vol. 104 | 23rd International Conference on Types for Proofs and Programs (TYPES 2017)
https://www.dagstuhl.de/dagpub/978-3-95977-071-2

Vol. 115 | 13th International Symposium on Parameterized and Exact Computation (IPEC 2018)
https://www.dagstuhl.de/dagpub/978-3-95977-084-2

Vol. 124 | 10th Innovations in Theoretical Computer Science Conference (ITCS 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-095-8

Vol. 125 | 22nd International Conference on Principles of Distributed Systems (OPODIS 2018)
https://www.dagstuhl.de/dagpub/978-3-95977-098-9

Vol. 126 | 36th International Symposium on Theoretical Aspects of Computer Science (STACS 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-100-9

Vol. 127 | 22nd International Conference on Database Theory (ICDT 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-101-6

Vol. 128 | 30th Annual Symposium on Combinatorial Pattern Matching (CPM 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-103-0

Vol. 129 | 35th International Symposium on Computational Geometry (SoCG 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-104-7

Vol. 130 | 24th International Conference on Types for Proofs and Programs (TYPES 2018)
https://www.dagstuhl.de/dagpub/978-3-95977-106-1

Vol. 131 | 4th International Conference on Formal Structures for Computation and Deduction (FSCD 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-107-8

Vol. 132 | 46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-109-2

Vol. 133 | 31st Euromicro Conference on Real-Time Systems (ECRTS 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-110-8

Vol. 134 | 33rd European Conference on Object-Oriented Programming (ECOOP 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-111-5

Vol. 135 | 14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-112-2

Vol. 136 | 3rd Summit on Advances in Programming Languages (SNAPL 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-113-9

Vol. 137 | 34th Computational Complexity Conference (CCC 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-116-0

Vol. 138 | 44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-117-7

Vol. 139 | 8th Conference on Algebra and Coalgebra in Computer Science (CALCO 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-120-7

Vol. 140 | 30th International Conference on Concurrency Theory (CONCUR 2019)
https://www.dagstuhl.de/dagpub/978-3-95977-121-4

Fig. 4.5
LIPIcs volumes published in 2019 – Part 1.

Vol. 141 | 10th International Conference on Interactive Theorem Proving (ITP 2019) https://www.dagstuhl.de/dagpub/978-3-95977-122-1
Vol. 142 | 14th International Conference on Spatial Information Theory (COSIT 2019) https://www.dagstuhl.de/dagpub/978-3-95977-115-3
Vol. 143 | 19th International Workshop on Algorithms in Bioinformatics (WABI 2019) https://www.dagstuhl.de/dagpub/978-3-95977-123-8
Vol. 144 | 27th Annual European Symposium on Algorithms (ESA 2019) https://www.dagstuhl.de/dagpub/978-3-95977-124-5
Vol. 145 | Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019) https://www.dagstuhl.de/dagpub/978-3-95977-125-2
Vol. 146 | 33rd International Symposium on Distributed Computing (DISC 2019) https://www.dagstuhl.de/dagpub/978-3-95977-126-9
Vol. 147 | 26th International Symposium on Temporal Representation and Reasoning (TIME 2019) https://www.dagstuhl.de/dagpub/978-3-95977-127-6
Vol. 148 | 14th International Symposium on Parameterized and Exact Computation (IPEC 2019) https://www.dagstuhl.de/dagpub/978-3-95977-129-0
Vol. 149 | 30th International Symposium on Algorithms and Computation (ISAAC 2019) https://www.dagstuhl.de/dagpub/978-3-95977-130-6
Vol. 150 | 39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2019) https://www.dagstuhl.de/dagpub/978-3-95977-131-3

Fig. 4.6
LIPIcs volumes published in 2019 – Part 2.

APPROX | International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX) accepted for 2019–2023 (Re-evaluation)
CCC | Computational Complexity Conference accepted for 2020–2024 (Re-evaluation)
CONCUR | International Conference on Concurrency Theory (CONCUR) accepted for 2020–2024 (Re-evaluation)
DNA | International Conference on DNA Computing and Molecular Programming accepted for 2020–2024
ECOOP | European Conference on Object-Oriented Programming accepted for 2020–2024 (Re-evaluation)
FORC | Symposium on Foundations of Responsible Computing accepted for 2020–2022
FSCD | Formal Structures for Computation and Deduction accepted for 2021–2025 (Re-evaluation)
ICDT | International Conference on Database Theory accepted for 2020–2024 (Re-evaluation)
ITC | Conference on Information-Theoretic Cryptography accepted for 2020–2022
ITP | International Conference on Interactive Theorem Proving (ITP) accepted for 2019–2023
RANDOM | International Conference on Randomization and Computation (RANDOM) accepted for 2019–2023 (Re-evaluation)
SoCG | Symposium on Computational Geometry accepted for 2020–2024 (Re-evaluation)

Fig. 4.7
Conferences accepted in 2019 for publication in LIPIcs.

Dagstuhl Publishing

LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*[30] veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift gegründet und in 2013 wurde der Betrieb aufgenommen. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 4.8), begutachtet alle eingereichten Arbeiten. Die Zeitschrift wird gemeinsam mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*[31] der Fachgesellschaft *European Design and Automation Association (EDAA)*[32] 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 2019 wurde eine Ausgabe von *LITES* mit insgesamt 5 Artikeln veröffentlicht.

LITES: Leibniz Transactions on Embedded Systems

The *LITES*[30] journal publishes original peer-reviewed articles on all aspects of embedded computer systems via Open Access. The journal was established in 2012 and started operating in early 2013. A broad team of experienced researchers, acting as editorial board (see Fig. 4.8), reviews all submitted contributions. The journal is jointly published with the *EMbedded Systems Special Interest Group (EMSIG)*[31] of the *European Design and Automation Association (EDAA)*[32]. 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 2019, one issue of *LITES* containing 5 articles in total was published.

[30] https://www.dagstuhl.de/lites
[31] http://www.emsig.net/
[32] https://www.edaa.com/

Fig. 4.8
LITES Editorial Board.

Dagstuhl Publishing

Infrastruktur 4.2 *Infrastructure*

Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 4.9. Die Bände aus den Reihen *LIPIcs* und *OASIcs* werden zudem bei Scopus[33] eingereicht, wo sie regelmäßig indiziert werden. Die Reihen *LIPIcs* und *OASIcs* sowie die Zeitschrift *LITES* sind zudem im Directory of Open Access Journals (DOAJ) gelistet, siehe Fig. 4.9. Zudem unterstützen die technischen Schnittstellen die Datenakquise durch GoogleScholar, so dass die Publikationen sichtbar und besser recherchierbar sind.

LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*[34] ein Online-Repositorium ins Leben gerufen, um Open Access-Veröffentlichungen von Leibniz-Instituten und deren Wissenschaftlern zu unterstützen und sichtbar zu machen. Schloss Dagstuhl liefert alle Artikel aus den Reihen *Dagstuhl Reports* und *Dagstuhl Manifestos* an das Repositorium und stärkt dadurch Forschungsergebnisse aus der Informatik innerhalb dieses multidisziplininären Repositoriums.

Indexing

All series of the publication portfolio are listed in *dblp*; see Fig. 4.9. The LIPIcs and OASIcs volumes are submitted to Scopus[33] where they are regularly indexed. The *LIPIcs* and *OASIcs* series as well as the journal *LITES* are also listed in the Directory of Open Access Journals (DOAJ), see Fig. 4.9. The technical interface of our publication server enables harvesting according to the Google Scholar guidelines. Google Scholar regularly retrieves metadata and full-texts from our server.

LeibnizOpen

The Leibniz Association has established the *LeibnizOpen*[34] 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.

[33] https://www.scopus.com
[34] http://www.leibnizopen.de/

dblp
Dagstuhl Reports https://dblp.org/db/journals/dagstuhl-reports/
Dagstuhl Manifestos https://dblp.org/db/journals/dagstuhl-manifestos/
Dagstuhl Follow-Ups https://dblp.org/db/series/dfu/
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. 4.9
Indexing of Dagstuhl Publishing series in dblp and DOAJ.

Dagstuhl Publishing

AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Predatory Publishing“[35] mit organisiert, welcher bereits der fünfte Workshop in Folge seit 2013 ist. Der Workshop fand am 17. und 18. Januar 2019 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt.

Publikationsserver: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)[36] werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative[37] berücksichtigt, wodurch alle nötigen Metadaten zu jeder Publikation gespeichert werden und die Langzeitverfügbarkeit sichergestellt wird. Die Online-Publikationen sind zitierfähig und stehen einer großen Leserschaft zur Verfügung.

Einreichungssystem: DSub

Im Frühjahr 2019 wurde ein von Dagstuhl entwickeltes Einreichungssystem names DSub eingeführt. Mit diesem System werden seit dem alle Einreichungen zu den Reihen LIPIcs und OASIcs entgegengenommen. Unter anderem wurde mit dem neuen System dem Wunsch einer aktiven Autorenfreigabe der überarbeiteten Dokumente vor der Veröffentlichung entsprochen und die automatische Extraktion der Metadaten aus den LaTeX-Quellen ermöglicht.

Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)[38] 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.[39]
- SunSite Central Europe: Der Sun-Server-Park, der an der RWTH Aachen betrieben wird, bietet eine Heimat für zahlreiche Software-Archive und Publikationen. Der gesamte DROPS-Bestand wird in regelmäßigen Abständen auf der SunSite Aachen gespiegelt.[40]

Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Predatory Publishing”[35] 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 17 and 18, 2019.

Publication server: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)[36]. The general guidelines of the Dublin Core initiative[37] 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.

Submission system: DSub

In spring 2019 a submission system called DSub developed by Dagstuhl was introduced. Since then, this system has been used to process all submissions for the LIPIcs and OASIcs series. Among other things, the new system has satisfied the need for active author approval of revised documents prior the publication and enables automatic extraction of metadata from LaTeX sources.

Long-term Archiving

All publications are submitted to the German National Library (D-NB)[38] 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[39].
- SunSite Central Europe: The Sun server park, located at the Aachen University of Technology, is home to numerous software archives and publications. All the DROPS assets are mirrored at regular intervals on the Aachen SunSite.[40]

[35] https://www.dagstuhl.de/fileadmin/dagpub/journalmanagement-leibniz/2019-01-workshop/
[36] https://www.dagstuhl.de/drops
[37] http://dublincore.org/
[38] https://www.dnb.de/DE/Professionell/Erhalten/erhalten_node.html#sprg209698
[39] https://subs.emis.de/LIPIcs/
[40] 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.

Charles Hansen, SCI Institute, University of Utah, on "Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization" having been downloaded 56407 times since 2014

11231 – Scientific Visualization | Dagstuhl Seminar | https://www.dagstuhl.de/11231

[...]the book generated from the successful Dagstuhl Seminar meeting 11231, Scientific Visualization: Uncertainty, Multifield, Biomedical, and Scalable Visualization, has been in the top 25% of downloads for Springer this past year (still, after all these years!) I hope these statistics help the Dagstuhl scientific mission, and indicates the value of having a large seminar at the Dagstuhl site.

Gary Burnett, in his Abstract for the Dagstuhl Report on 19132 - Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles?

16262 – Automotive User Interfaces in the Age of Automation | Dagstuhl Seminar | https://www.dagstuhl.de/16262

As for 2016 when I was here, I have benefited immensely from the in-depth and extended conversations with colleagues on human-centered design issues for future vehicles. I have learnt about many interesting theories, methods, studies that can inform our work and look forward to working closely with my Dagstuhl friends in years to come

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.

Julien Cornebise

19082 – AI for the Social Good | Dagstuhl Seminar | https://www.dagstuhl.de/19082

Many thanks Dagstuhl team for sharing the survey results – in addition to the invaluable hosting of this seminar.

Britta Dorn

19443 – Algorithms and Complexity in Phylogenetics | Dagstuhl Seminar | https://www.dagstuhl.de/19443

Ganz besonders möchte ich mich auch beim Küchenteam bedanken, das so freundlich und liebevoll auf alle unsere Sonderwünsche eingegangen ist und tatsächlich jedes Mal etwas gefunden hat, das unsere Kinder gerne essen wollten! Meine Familie hat sich im Schloss ebenfalls sehr wohlgefühlt und ich bin richtig glücklich, dass ich die Kinder mitnehmen konnte. Vielen Dank dafür!

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 19052

19052 – Computational Methods for Melody and Voice Processing in Music Recordings | Dagstuhl Seminar | https://www.dagstuhl.de/19052

Many of our participants were visiting Dagstuhl for the first time and enthusiastically praised the open and inspiring setting. The group dynamics were excellent with many personal exchanges and common activities. Some scientists expressed their appreciation for having the opportunity for prolonged discussions with researchers from neighboring research fields—something that is often impossible during conference-like events. In conclusion, our expectations for the seminar were not only met but exceeded, in particular concerning networking and community building. We want to express our gratitude to the Dagstuhl board for giving us the opportunity to organize this seminar, the Dagstuhl office for their exceptional support in the organization process, and the entire Dagstuhl staff for their excellent service during the seminar.

Organizers of Dagstuhl Seminar 19132

19132 – Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles? | Dagstuhl Seminar | https://www.dagstuhl.de/19132

[...]as organizers, we would like to express our deep appreciation to all of those people who contributed to the success of this workshop. First and foremost, we thank the team at Schloss Dagstuhl for their dedication and exceptionally high-quality work, from organizing the meeting, to hosting us at the castle.[...]

Organizers of Dagstuhl Seminar 19021

19021 – Joint Processing of Language and Visual Data for Better Automated Understanding | Dagstuhl Seminar | https://www.dagstuhl.de/19021

The organizers would like to thank the Dagstuhl team for their continuous support; the welcoming atmosphere made the seminar both highly productive and enjoyable.

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.

Sergi Valverde (UPF - Barcelona, ES)

Twitter | https://twitter.com/svalver/status/1206900285261660160

Wow! Dagstuhl is one of the best places in Europe to hold working groups and meetings. Many thanks for sharing your notes!

Johanna Pirker (TU Graz, AT)

19272 – Real VR - Importing the Real World into Immersive VR and Optimizing the Perceptual Experience of Head-Mounted Displays | Dagstuhl Seminar | https://twitter.com/JoeyPrink/status/1145451766730907654

I arrived at my very first @dagstuhl seminar "Real VR".. this is just such a fantastic place!!! with a piano, lots of nature and sports, a wine cellar AND a 24/7 CS library next to my room!!!!! I am home. #VR #ar #dagstuhl

Neil Ernst

Twitter | https://twitter.com/neilernst/status/1129161168625520640

Crazy acknowledgement in this ICSE paper: “The first and third authors are deeply grateful to the organizers of Dagstuhl Seminar 07491, without whose kind invitations we likely would have never met, married, or written this paper.” Presumably the first 2 are the important ones ..

Cindy L. Bethel (Mississippi State University, US)

19411 – Social Agents for Teamwork and Group Interactions | Dagstuhl Seminar | https://twitter.com/bethelcl/status/1106489005212749824

I am super excited and honored to be invited to participate in one of the Dagstuhl Seminars in Germany this fall that are quite prestigious internationally in Computer Science! I am really looking forward to this week-long sharing and discussions with colleagues around the world!

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.

19021 – Joint Processing of Language and Visual Data for Better Automated Understanding | Dagstuhl Seminar | https://www.dagstuhl.de/19021

It’s great that Dagstuhl offers this type of seminars, which are always a source of inspiration and a perfect getaway to focus on longer term research directions.

19021 – Joint Processing of Language and Visual Data for Better Automated Understanding | Dagstuhl Seminar | https://www.dagstuhl.de/19021

It was valuable to be able to look up papers and publications in paywalled journals, thank you.

19032 – Conditional Logics and Conditional Reasoning: New Joint Perspectives | Dagstuhl Seminar | https://www.dagstuhl.de/19032

It is always a wonderful academic experience to be at Schloss Dagstuhl. Thank you very much!

19072 – The Role of Non-Monotonic Reasoning in Future Development of Artificial Intelligence | Dagstuhl Seminar | https://www.dagstuhl.de/19072

Dagstuhl is really unique and perfect.

19081 – Verification and Synthesis of Human-Robot Interaction | Dagstuhl Seminar | https://www.dagstuhl.de/19081

It was overall an amazing experience. It makes me wish the rest of my professional life worked this way with time to think and deeply engaged colleagues to work with. Thank you!

19082 – AI for the Social Good | Dagstuhl Seminar | https://www.dagstuhl.de/19082

All the people at Dagstuhl were lovely - what an amazing group of people you have working at the center.

19101 – Analysis, Design, and Control of Predictable Interconnected Systems | Dagstuhl Seminar | https://www.dagstuhl.de/19101

As always, Dagstuhl is an excellent venue with superb organization.

19101 – Analysis, Design, and Control of Predictable Interconnected Systems | Dagstuhl Seminar | https://www.dagstuhl.de/19101

The Dagstuhl concept continues being a reference in the computer science community, world wide. Congratulations.

19101 – Analysis, Design, and Control of Predictable Interconnected Systems | Dagstuhl Seminar | https://www.dagstuhl.de/19101

Thank you for your excellent format. You are doing a great service to the academic community. The impact is immense. It would be great if you could extend this to other countries. For instance, can we think of a similar format in India. I am sure it will be very welcome and immensely valuable, though operational challenges may exist.

19111 – Theoretical Foundations of Storage Systems | Dagstuhl Seminar | https://www.dagstuhl.de/19111

This was a really amazing experience. Much better than a conference. I think the things that made it great were: environment and setting; particularly I would like to commend the quality of the food and coffee, and the library (which is simply amazing, and on its own would be enough to justify a trip) people - the quality of the people, and the mix from different areas (industry, academia, theory, applications) timescale - sufficient to allow attendees to get to know each other beyond the superficial, which is really really important for collaboration quality of research on display

19112 – Engineering Reliable Multiagent Systems | Dagstuhl Seminar | https://www.dagstuhl.de/19112

I especially like the cheese in the evening, it brings people together.

19132 – Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles? | Dagstuhl Seminar | https://www.dagstuhl.de/19132

This was perhaps the highest density of academically rewarding activity of the things I have for the year.

19132 – Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles? | Dagstuhl Seminar | https://www.dagstuhl.de/19132

I think Schloss Dagstuhl is at the cutting edge of international research collaboration support activities.

19141 – Programmable Network Data Planes | Dagstuhl Seminar | https://www.dagstuhl.de/19141

Dagstuhl is unique in the world - your efforts are greatly appreciated and it is an honor to be invited here.

19141 – Programmable Network Data Planes | Dagstuhl Seminar | https://www.dagstuhl.de/19141

i think i've just seen things getting better along all dimensions

19141 – Programmable Network Data Planes | Dagstuhl Seminar | https://www.dagstuhl.de/19141

More and more conferences and other meetings try to arrange childcare and struggle with it. At Dagstuhl, not only is it available, it is also by professional caretakers, affordable and well organized. As I said above, PLEASE advertise this MUCH more prominently - it helps significantly in attracting younger and esp. female participants!

19181 – Computational Geometry | Dagstuhl Seminar | https://www.dagstuhl.de/19181

The best aspect is definitely the fact that putting researchers together in an isolated environment gives plenty opportunities for them to work together and exchange ideas, much more so than any conferences.

19181 – Computational Geometry | Dagstuhl Seminar | https://www.dagstuhl.de/19181

I want to thank the wonderful staff who prepared the meals and served in such a nice way. I came here with my family, all of which prefers to eat vegetarian, and all of us were so happy with the meals. The staff were so friendly towards kids, and it made our stay really comfortable. Thank you!

19181 – Computational Geometry | Dagstuhl Seminar | https://www.dagstuhl.de/19181

Best facility for scientific research I have ever seen.

19181 – Computational Geometry | Dagstuhl Seminar | https://www.dagstuhl.de/19181

If it wouldn't be for the family room without extra costs, and the extra family friendly environment of Dagstuhl; I wouldn't be able to attend this seminar. I want to thank all who made my visit possible.

19211 – Enumeration in Data Management | Dagstuhl Seminar | https://www.dagstuhl.de/19211

The format of Dagstuhl seminars is really great. The intensive interaction with colleagues from my own field _and_ from related fields could not be better than here.

19211 – Enumeration in Data Management | Dagstuhl Seminar | https://www.dagstuhl.de/19211

Change nothing! This is perfect.

19211 – Enumeration in Data Management | Dagstuhl Seminar | https://www.dagstuhl.de/19211

Dagstuhl is a very special resource for computer science research. I try to come every time I am invited, the long trip notwithstanding. Thank you.

19232 – Ubiquitous Computing Education: Why, What, and How | Dagstuhl Seminar | https://www.dagstuhl.de/19232

This is a fantastic place. Either the staff read in researchers mind, or they listen carefully to their concerns, practices and demands. I wish there was more places like this around the world.

19232 – Ubiquitous Computing Education: Why, What, and How | Dagstuhl Seminar | https://www.dagstuhl.de/19232

The meals were all nicely prepared and served. I appreciate the mixing of the name tags on different tables. It's a nice and thoughtful gesture. The rooms are very spacious and nicely appointed. The rooms are also quite accessible, so I appreciate that. The newspaper rooms, wine cellar, coffee room, and the music rooms are all very nicely arranged. It's a delight to wander around in the Castle and the new Building and the library too.

19232 – Ubiquitous Computing Education: Why, What, and How | Dagstuhl Seminar | https://www.dagstuhl.de/19232

Dagstuhl is a fantastic initiative with a fantastic venue. This is possibly the most useful academic activity in my calendar this year.

19232 – Ubiquitous Computing Education: Why, What, and How | Dagstuhl Seminar | https://www.dagstuhl.de/19232

Both in my previous and current experience, I deeply respect, value, and appreciate Dagstuhl. I know of no other organization in the world which either aspires to or realizes a similar role. While admittedly a heterogeneous comparison, I've increasingly made analogies between Dagstuhl and the BBC relative to intellectual impact on the world community (as someone whom has for decades held high regard and frequent engagement with the BBC). Many thanks to the organization and all of its sponsors and enablers for this invaluable world service!

19232 – Ubiquitous Computing Education: Why, What, and How | Dagstuhl Seminar | https://www.dagstuhl.de/19232

I am happy to learn about this incredible library with the large physical collection of publications. I am also delighted to know that the D in DBLP is Dagstuhl. Thank you.

19232 – Ubiquitous Computing Education: Why, What, and How | Dagstuhl Seminar | https://www.dagstuhl.de/19232

I don't think there is other place offering what Dagstuhl offers and I think it is really important for the community to keep Dagstuhl alive for long, or use it as an example to develop such places. I can see that people all over the world know what is Dagstuhl and want to experience it because all know this is not only a rich experience but also a fun and relaxing and sociale event. A great combination for moving a field forward.

19272 – Real VR - Importing the Real World into Immersive VR and Optimizing the Perceptual Experience of Head-Mounted Displays | Dagstuhl Seminar | https://www.dagstuhl.de/19272

I was a bit suspicious about the usefulness of making a trip to a remote place in Germany, but I am very glad that I attended the Dagstuhl seminar. I will definitely participate in the future if chances allow.

19281 – Notional Machines and Programming Language Semantics in Education | Dagstuhl Seminar | https://www.dagstuhl.de/19281

Thank you for inviting me. This was a very influential point for my career. The impact of Dagstuhl is not just on communities/disciplines. It also affects (positively) many individual people. This is important.

19281 – Notional Machines and Programming Language Semantics in Education | Dagstuhl Seminar | https://www.dagstuhl.de/19281

This seminar was one of the best research experiences I have had in my 20+ year career. Many factors contributed to its success for me including the topic and its timeliness in my career. However, most notable was how the program and the Dagstuhl traditions (e.g. randomly assigned seatings) was able to consistently draw from all of its participants, despite (or perhaps because of) the diversity.

19291 – Values in Computing | Dagstuhl Seminar | https://www.dagstuhl.de/19291

The Dagstuhl infrastructure is great. I especially liked that lunch and dinner were served on-site (much time to talk) and that the participants were mixed at the tables. This meant that you met many new people during the meals and the random seating arrangement is a great mixer and a great equaliser.

19292 – Mobile Data Visualization | Dagstuhl Seminar | https://www.dagstuhl.de/19292

Awesome possibility for doing sport activities in groups, self-service drinks & coffee are awesome.

19292 – Mobile Data Visualization | Dagstuhl Seminar | https://www.dagstuhl.de/19292

Vegan food was amazing! Big thanks to the chef!

19292 – Mobile Data Visualization | Dagstuhl Seminar | https://www.dagstuhl.de/19292

The library is exceptional, including personnel & physical connection. I found an article that I had been looking for forever! Much appreciated.

19292 – Mobile Data Visualization | Dagstuhl Seminar | https://www.dagstuhl.de/19292

I like having all the books from seminar participants there!

19301 – Secure Composition for Hardware Systems | Dagstuhl Seminar | https://www.dagstuhl.de/19301

I was surprised and happy to see that it is possible to have a free access to leading journals in my area. It helped a lot when we wanted to discuss things in depth. Thank you!

19331 – Software Protection Decision Support and Evaluation Methodologies | Dagstuhl Seminar | https://www.dagstuhl.de/19331

The location and quietness of the venue lets focus entirely on the topic. The self-service and fair use policy makes the whole experience as fluid as possible.

19331 – Software Protection Decision Support and Evaluation Methodologies | Dagstuhl Seminar | https://www.dagstuhl.de/19331

Keep Dagstuhl as it is. No tweets, no bullshit, just science.

19341 – Algorithms and Complexity for Continuous Problems | Dagstuhl Seminar | https://www.dagstuhl.de/19341

You just fell like coming home.

19341 – Algorithms and Complexity for Continuous Problems | Dagstuhl Seminar | https://www.dagstuhl.de/19341

Dagstuhl creates such a perfect environment that you can fully concentrate on research and people. From my own example and many others, Dagstuhl provides incredible opportunities to young researchers!

19352 – Computation in Low-Dimensional Geometry and Topology | Dagstuhl Seminar | https://www.dagstuhl.de/19352

Everything is great, including the music room. Just keep it as it is! The fact that families are welcome is definitely a plus.

19361 – Logic and Learning | Dagstuhl Seminar | https://www.dagstuhl.de/19361

Very friendly environment. Everything is designed to create a small family/community atmosphere. I like that the castle is so remote.

19381 – Application-Oriented Computational Social Choice | Dagstuhl Seminar | https://www.dagstuhl.de/19381

Just keep 'being Dagstuhl'! You are setting an example on how to foster true science worldwide.

19401 – Comparative Theory for Graph Polynomials | Dagstuhl Seminar | https://www.dagstuhl.de/19401

the facilities really encourage collaboration and collegiality, and were well-used, with high-level research questions richly discussed in the evenings.

19411 – Social Agents for Teamwork and Group Interactions | Dagstuhl Seminar | https://www.dagstuhl.de/19411

Special dietary requirements were really taken well care of, really nice.

19421 – Quantum Cryptanalysis | Dagstuhl Seminar | https://www.dagstuhl.de/19421

The atmosphere is perfect for researchers to discuss relevant problems and directions for future study. No distractions, and plenty of available resources.

19431 – Theory of Randomized Optimization Heuristics | Dagstuhl Seminar | https://www.dagstuhl.de/19431

I started two research collaborations this week. The previous Dagstuhl seminar had a very big impact on my personal research agenda since then, and I think the same has happened again. I got a lot of ideas for future research.

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.

zbMATH (Berlin, Germany)

Twitter | https://twitter.com/zbMATH/status/1081221764678631424

Thanks to our partners from @dagstuhl/@dblp_org and @HITStudies for three years of great collaboration in the @LeibnizWGL project Scalable Author Disambiguation for Bibliographic Databases. Implementation of the results to #zbMATH has already yielded connection with many new sources of author information (almost 10 times more external identifiers interlinked with #zbMATH author profiles), >10 times (mostly automatic) exclusions of publications from authors (despite homonymous signatures), and a significant decrease of ambiguous assignments.

Gabor Szarnyas (Budapest University of Technology and Economics, Hungary)

Twitter | https://twitter.com/szarnyasg/status/1084709650619133952

The number of publications on @dblp_org just hit 4,444,444. Keep going and thanks for your work.

Vijay Chidambaram (University of Texas at Austin, TX, USA)

Twitter | https://twitter.com/vj_chidambaram/status/1096059540238409732

What really gets me is that @dblp_org is a free service, and they are already so good! Imagine what they could do with some part of 21 million!

Nate Foster (Cornell University, NY)

Twitter | https://twitter.com/natefoster/status/1096057539718979584

Since it's Valentine's Day: I ♡ @dblp_org

Olivier Jacquot (Paris, France)

Twitter | https://twitter.com/olivejacquot/status/1106210963987025921

Tout acteur des Humanités numériques doit connaître (si ce n'est pas déjà le cas) le travail fait par l'université de Trier et le Schloss Dagstuhl - Leibniz-Zentrum für Informatik GmbH (LZI): https://dblp.uni-trier.de[41]

Alireza Ghasemi (Zurich, Switzerland)

Twitter | https://twitter.com/a_ghasemi/status/1112457686967377925

For some not-easily-explainable reason, I have always preferred @dblp_org to #GoogleScholar. For me, it has always yielded far less false positives, hence making it way faster to get yo the results, especially when you know what you are looking for.

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

Twitter | https://twitter.com/mm_jj_nn/status/1138299362734481408

A thing I like about DBLP is that, though also far from perfect, they're willing to accept that this isn't an automatable problem at the moment, and regularly do manual corrections (I've emailed in 50 corrections, and all were fixed within a week or two). Feels more responsible.

[41] engl.: Anyone involved in the Digital Humanities should be aware (if they are not already) of the work done by the University of Trier and the Schloss Dagstuhl - Leibniz Center for Informatics GmbH (LZI): https://dblp.uni-trier.de

Danica Radovanović (Oxford, UK)

Twitter | https://twitter.com/danicar/status/1171423071359516672

I'm a Social Scientist and I'm humbled to have some of my bibliography items in Computer Science Bibliography Database.

Thomas Vogel (Humboldt University Berlin, Germany)

Twitter | https://twitter.com/tomvog/status/1194668680765947906

This is pretty cool: @dblp_org has integrated @unpaywall! Now links to open-access versions of papers are shown in dblp. #openaccess

John Samuel (CPE Lyon, France)

Twitter | https://twitter.com/jsamwrites/status/1194696753490083841

Great to see #OpenCitations data being integrated to DBLP, now that more than 50% of citations are openly available.

Another major development coming from @dblp_org after #OpenCitations. The data dumps are now released under CC0 licence. Indeed a major step towards open science movement. #opendata

Nick Walker (University of Washington, Seattle, WA, USA)

Twitter | https://twitter.com/nickwalker_us/status/1194848350173949954

Also, kudos to @dblp_org for disambiguating my entry within 24 hours of asking. You're awesome, and the source of my expectations for how to do author profiles well!

Jan Martin Keil (University of Jena, Germany)

Twitter | https://twitter.com/janmartinkeil/status/1194724449565773826

@dblp_org becomes even more helpful by integration of citation data. Thank you.

Victoria X. Lin (Palo Alto, CA, USA)

Twitter | https://twitter.com/VictoriaLinML/status/1195623532375330816

I somehow find the Bibtex format in DBLP better than G-Scholar and many other automatic tools. This is the inspiration for this tool. #AcademicLife

A.C. (@actsant)

Twitter | https://twitter.com/actsant/status/1196897704128987137

Feliz por descobrir que tenho uma página no dblp. Só tem um artigo lá, mas já é um começo.[42]

[42] engl.: Happy to find out I have a page on dblp. There's only one article there, but it's a start.

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.

Roberto Di Cosmo

Twitter | https://twitter.com/rdicosmo/status/1135902609020264448

I'm honored and thrilled to join the editorial board of #LIPIcs, a series of high quality #Open #Access conference proceedings from computer science conferences worldwide, and an important contribution to #Open #Science https://dagstuhl.de/en/publications/lipics/editorial-board/… @inria @SWHeritage @IRIF_Paris

Michael Greenberg

Twitter | https://twitter.com/mgrnbrg/status/1090258658460684288

@dagstuhl is a treasure—not just for its fun seminars, but for LIPIcs (bonus points for just putting the bibtex on the page)

Editor

Author/Editor Survey | https://www.dagstuhl.de/

Excellent work, I'm very happy with everything! From the perspective of the PC chair, the new process is really smooth, professional, and easy to use.

Author

Author/Editor Survey | https://www.dagstuhl.de/

The service level is fantastic. In particular we are so thankful you allowed us to correct our own minor errors (typos, minor miswordings, two wrong references) in our final submitted version. Thanks so much!

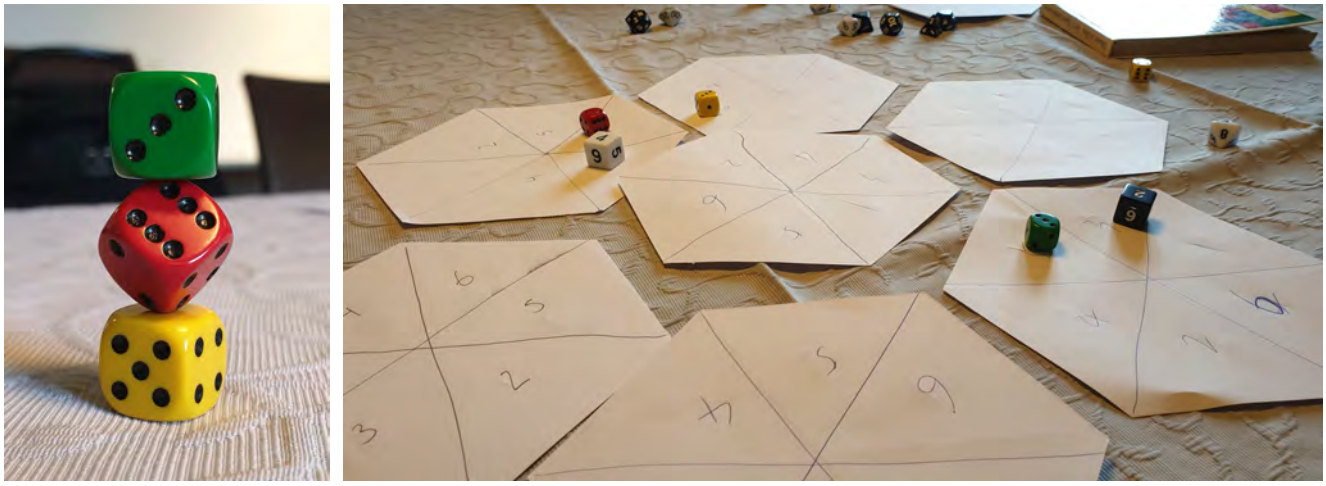

Fig. 5.1
„Since then the work group I've joined has focused on 'impossible games'. Games that based on elements such as cognitive load, social evaluation or minute physical interactions, becomes impossible for an AI – or a human – to play effectively." Twitter post by 19511 Dagstuhl Seminar participant Tommy Thompson. https://twitter.com/AIandGames/status/1207280884153552896. Photo courtesy of Tommy Thompson.

6 Die Seminare in 2019
The 2019 Seminars

Applications, Interdisciplinary Work

- Advances and Challenges in Protein-RNA Recognition, Regulation and Prediction (19342)
- Computational Creativity Meets Digital Literary Studies (19172)
- Computational Proteomics (19351)
- Computational Methods for Melody and Voice Processing in Music Recordings (19052)
- Ethics and Trust: Principles, Verification and Validation (19171)
- Future Automotive HW/SW Platform Design (19502)
- Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles? (19132)
- Values in Computing (19291)

Artificial Intelligence, Computational Linguistics

- AI for the Social Good (19082)
- Analysis of Autonomous Mobile Collectives in Complex Physical Environments (19432)
- Application-Oriented Computational Social Choice (19381)
- Artificial and Computational Intelligence in Games: Revolutions in Computational Game AI (19511)
- Conversational Search (19461)
- Diversity, Fairness, and Data-Driven Personalization in (News) Recommender System (19482)
- Emerging Hardware Techniques and EDA Methodologies for Neuromorphic Computing (19152)
- Joint Processing of Language and Visual Data for Better Automated Understanding (19021)
- Social Agents for Teamwork and Group Interactions (19411)
- The Role of Non-Monotonic Reasoning in Future Development of Artificial Intelligence (19072)
- Ubiquitous Computing Education: Why, What, and How (19232)

Cryptography, Security, Privacy

- Biggest Failures in Security (19451)
- Cybersafety Threats – from Deception to Aggression (19302)
- Distributed Computing with Permissioned Blockchains and Databases (19261)
- Empirical Evaluation of Secure Development Processes (19231)
- Practical Yet Composably Secure Cryptographic Protocols (19042)
- Quantum Cryptanalysis (19421)
- Secure Composition for Hardware Systems (19301)
- Software Protection Decision Support and Evaluation Methodologies (19331)

Databases, Information Retrieval, Machine Learning, Data Mining

- 25 Years of the Burrows-Wheeler Transform (19241)
- Big Graph Processing Systems (19491)
- Data Ecosystems: Sovereign Data Exchange among Organizations (19391)
- Data Series Management (19282)
- Logic and Learning (19361)
- Machine Learning Meets Visualization to Make Artificial Intelligence Interpretable (19452)
- Multi-Document Information Consolidation (19182)

Data Structures, Algorithms, Complexity

- Algorithmic Problems in Group Theory (19131)
- Algorithms and Complexity in Phylogenetics (19443)
- Algorithms and Complexity for Continuous Problems (19341)
- Beyond-Planar Graphs: Combinatorics, Models and Algorithms (19092)
- Comparative Theory for Graph Polynomials (19401)
- Computation in Low-Dimensional Geometry and Topology (19352)
- Computational Complexity of Discrete Problems (19121)
- Computational Geometry (19181)
- Data Structures for the Cloud and External Memory Data (19051)
- Enumeration in Data Management (19211)
- Graph Colouring: from Structure to Algorithms (19271)
- New Horizons in Parameterized Complexity (19041)
- Theoretical Foundations of Storage Systems (19111)
- Theory of Randomized Optimization Heuristics (19431)
- Topology, Computation and Data Analysis (19212)

Distributed Computation, Networks, Architecture, Systems

- Analysis, Design, and Control of Predictable Interconnected Systems (19101)
- Control of Networked Cyber-Physical Systems (19222)
- Programmable Network Data Planes (19141)
- Programming Languages for Distributed Systems and Distributed Data Management (19442)

Geometry, Image Processing, Graphics, Visualization

- 3D Morphable Models (19102)
- Astrographics: Interactive Data-Driven Journeys through Space (19262)
- Interactive Design and Simulation (19512)
- Mobile Data Visualization (19292)
- Real VR – Importing the Real World into Immersive VR and Optimizing the Perceptual Experience of Head-Mounted Displays (19272)
- Visual Analytics for Sets over Time and Space (19192)
- Visual Analytics of Multilayer Networks Across Disciplines (19061)
- Visual Computing in Materials Sciences (19151)

Software Technology, Programming Languages

- Approaches and Applications of Inductive Programming (19202)
- BOTse: Bots in Software Engineering (19471)
- Notional Machines and Programming Language Semantics in Education (19281)
- Software Evolution in Time and Space: Unifying Version and Variability Management (19191)

Verification, Logic, Formal Methods, Semantics

- Bringing CP, SAT and SMT together: Next Challenges in Constraint Solving (19062)
- Composing Model-Based Analysis Tools (19481)
- Conditional Logics and Conditional Reasoning: New Joint Perspectives (19032)
- Deduction Beyond Satisfiability (19371)
- Engineering Reliable Multiagent Systems (19112)
- Logics for Dependence and Independence (19031)
- Specification Formalisms for Modern Cyber-Physical Systems (19071)
- Verification and Synthesis of Human-Robot Interaction (19081)

6.1 Joint Processing of Language and Visual Data for Better Automated Understanding

Organizers: Marie-Francine Moens, Lucia Specia, and Tinne Tuytelaars

Seminar No. 19021

Date: January 6–11, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.1.1

Participants: Zeynep Akata, Andrei Barbu, Loïc Barrault, Raffaela Bernardi, Thales Bertaglia, Ozan Caglayan, Stephen Clark, Luísa Coheur, Guillem Collell, Vera Demberg, Desmond Elliott, Aykut Erdem, Erkut Erdem, Raquel Fernández, Orhan Firat, Anette Frank, Stella Frank, Lisa Anne Hendricks, David C. Hogg, Frank Keller, Douwe Kiela, Dietrich Klakow, Chiraag Lala, Marius Leordeanu, Jindrich Libovický, Pranava Madhyastha, Florian Metze, Marie-Francine Moens, Siddharth Narayanaswamy, Jean Oh, Pavel Pecina, Bernt Schiele, Carina Silberer, Lucia Specia, Tinne Tuytelaars, Jakob Verbeek, David Vernon, Josiah Wang

The joint processing of language and visual data has recently received a lot of attention. This emerging research field is stimulated by the active development of deep learning algorithms. For instance, deep neural networks (DNNs) offer numerous opportunities to learn mappings between the visual and language media and to learn multimodal representations of content. Furthermore, deep learning recently has become a standard approach for automated image and video captioning and for visual question answering, the former referring to the automated description of images or video with descriptions in natural language sentences, the latter to the automated formulation of an answer in natural language to a question in natural language about an image.

Apart from aiding image understanding and the indexing and search of image and video data through the natural language descriptions, the field of jointly processing language and visual data builds algorithms for grounded language processing where the meaning of natural language is based on perception and/or actions in the world. Grounded language processing contributes to automated language understanding and machine translation of language. Recently, it has been shown that visual data provide world and common-sense knowledge that is needed in automated language understanding.

Joint processing of language and visual data is also interesting from a theoretical point of view for developing theories on the complementarity of such data in human(-machine) communication, for developing suitable algorithms for learning statistical knowledge representations informed by visual and language data, and for inferencing with these representations.

Given the current trend and results of multimodal (language and vision) research, it can be safely assumed that the joint processing of language and visual data will only gain in importance in the future. During the seminar we have discussed theories, methodologies and real-world technologies for joint processing of language and vision, particularly in the following research areas:

- Theories of integrated modelling and representation learning of language and vision for computer vision and natural language processing tasks;
- Explainability and interpretability of the learned representations;
- Fusion and inference based on visual, language and multimodal representations;
- Understanding human language and visual content;
- Generation of language and visual content;
- Relation to human learning;
- Datasets and tasks.

The discussions have attempted to give an answer to the following research questions (a non-exhaustive list):

- Which machine learning architectures will be best suited for the above tasks?
- How to learn multimodal representations that are relational and structured in nature to allow a structured understanding?
- How to generalize to allow recognitions that have few or zero examples in training?
- How to learn from limited paired data but exploiting monomodal models trained on visual or language data?
- How to explain the neural networks when they are trained for image or language understanding?
- How to disentangle the representations: factorization to separate the different factors of variation and discovering of their meaning?
- How to learn continuous representations that describe semantics and that integrate world and common-sense knowledge?
- How to reason with the continuous representations?
- How to translate to another modality?
- What would be effective novel evaluation metrics?

This Dagstuhl Seminar has brought together an interdisciplinary group of researchers from computer vision, natural language processing, machine learning and artificial intelligence to discuss the latest scientific realizations and to develop a roadmap and research agenda.

6.2 Logics for Dependence and Independence

Organizers: Erich Grädel, Phokion G. Kolaitis, Juha Kontinen, and Heribert Vollmer

Seminar No. 19031

Date: January 13–18, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.1.28

Participants: Leopoldo Bertossi, Dietmar Berwanger, Meghyn Bienvenu, Joachim Biskup, Katrin M. Dannert, Anuj Dawar, Arnaud Durand, Fredrik Engström, Bernd Finkbeiner, Floris Geerts, Erich Grädel, Gianluca Grilletti, Miika Hannula, Lauri Hella, Åsa Hirvonen, Matthias Hoelzel, Benny Kimelfeld, Phokion G. Kolaitis, Juha Kontinen, Paris Koutris, Sebastian Link, Martin Lück, Yasir Mahmood, Arne Meier, Magdalena Ortiz, Martin Otto, Eric J. Pacuit, Henri Prade, David J. Pym, Raine Rönnholm, Dan Suciu, Val Tannen, Bernhard Thalheim, Jouko Väänänen, Jan Van den Bussche, Jonni Virtema, Heribert Vollmer, Jef Wijsen, Richard Wilke, Fan Yang

Brief Introduction to the Topic

Dependence and independence are interdisciplinary notions that are pervasive in many areas of science. They appear in domains such as mathematics, computer science, statistics, quantum physics, and game theory. The development of logical and semantical structures for these notions provides an opportunity for a systematic approach, which can expose surprising connections between different areas, and may lead to useful general results.

Dependence Logic is a tool for modeling dependencies and interaction in dynamical scenarios. Reflecting this, it has higher expressive power and complexity than classical logics used for these purposes previously. Algorithmically, first-order dependence logic corresponds exactly to the complexity class NP and to the so-called existential fragment of second-order logic. Since the introduction of dependence logic in 2007, the framework has been generalized, e. g., to the contexts of modal, intuitionistic, and probabilistic logic. Moreover, interesting connections have been found to complexity theory, database theory, statistics, and dependence logic has been applied in areas such as linguistics, social choice theory, and physics. Although significant progress has been made in understanding the computational side of these formalisms, still many central questions remain unsolved so far. In addition to addressing the open questions, the seminar also aimed at boosting the exchange of ideas and techniques between dependence logic and its application areas.

Organization of the Seminar and Activities

The workshop brought together 40 researchers from mathematics, database theory, natural language semantics, and theoretical computer science. The participants consisted of both senior and junior researchers, including a number of postdocs and advanced graduate students.

Participants were invited to present their work and to communicate state-of-the-art advances. Over the five days of the workshop, 27 talks of various lengths took place. Introductory and tutorial talks of 90-60 minutes were scheduled prior to the workshop. Most of the remaining slots were filled, mostly with shorter talks, as the workshop commenced. The seminar ended with an open problems and perspectives session. The organizers considered it important to leave ample free time for discussion.

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

- Miika Hannula: Team semantics
- Val Tannen: Provenance
- Dan Suciu: Probabilistic databases
- Meghyn Bienvenu: Constraints in ontology based databases
- David Pym: Resource semantics
- Magdalena Ortiz: Complete and incomplete information in knowledge-enriched databases
- Jef Wijsen: Database repairs

In addition, the seminar consisted of 20 shorter contributed talks, addressing various topics concerning expressibility, axiomatizability, complexity and applications of team-based logics.

The last session of the workshop was devoted to open problems and consisted of contributions by Phokion Kolaitis, Jouko Väänänen and Juha Kontinen presenting questions about decidability and axiomatizability of the implication problem of various fragments of dependence and independence logic, Joachim Biskup addressing decidable first-order prefix classes in the database context, Heribert Vollmer presenting open rela-

tionships among various counting classes related to team-based logics, Lauri Hella talking about union-closed properties in Σ^1_1, and finally Raine Rönnholm addressing relationships between fragments of inclusion logic and greatest fixed-point logic.

The workshop ended with a discussion of future perspectives of the study of logics for dependence and independence.

The workshop achieved its aim of bringing together researchers from various related communities to share state-of-the-art research. The organizers left ample time outside of this schedule of talks and many fruitful discussions between participants took place throughout the afternoons and evenings.

Concluding Remarks and Future Plans

The organizers regard the workshop as a great success. Bringing together researchers from different areas fostered valuable interactions and led to fruitful discussions. Feedback from the participants was very positive as well.

Finally, the organizers wish to express their gratitude toward the Scientific Directorate of the Center for its support of this workshop.

6.3 Conditional Logics and Conditional Reasoning: New Joint Perspectives

Organizers: Guillaume Aucher, Paul Egré, Gabriele Kern-Isberner and Francesca Poggiolesi

Seminar No. 19032

Date: January 13–16, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.1.9.47

Participants: Guillaume Aucher, Giosué Baggio, Christoph Beierle, Didier Dubois, Giovanni Casini, Nicole Cruz de Echeverria Loebell, Paul Egré, Marianna Girlando, Giuseppe Greco, Mario Günther, Andreas Herzig, Andreas Kapsner, Stefan Kaufmann, Gabriele Kern-Isberner, Karolina Krzyzanowska, Nicola Olivetti, Francois Olivier, Graham Priest, Eric Raidl, Keith Stenning, Jakub Szymanik, Hans Rott, Niels Skovgaard Olsen, Ivan José Varzinczak, Emil Weydert

Logic in the first half of the 20th century has been mostly concerned with mathematical reasoning and providing a unified framework for the foundations of mathematics. In the second half of the 20th century, with the emergence of artificial intelligence, new formalisms have been introduced to model kinds of inference closer to everyday life.

"Commonsense reasoning", the reasoning that humans perform in everyday life, is significantly different from the reasoning of mathematicians, which has been the object of study of (mathematical) logic for a long time. It is very rich and includes different kinds of reasoning, such as counterfactual reasoning, default reasoning or uncertain and plausible reasoning. Commonsense reasoning is often captured by means of conditionals, which are sentences of the form 'if A then B'. These conditionals can also be of various kinds: counterfactual, indicative, or subjunctive. The benefits of conditionals for formalizing commonsense reasoning are basically twofold: first, they can encode reasoning patterns of various types if one chooses suitable semantics or calculi, and second, they provide a common syntactic element that can be used to relate and compare the different kinds of commonsense reasoning as well as the mathematical reasoning.

Conditionals are also studied in the psychology of reasoning, which has recently witnessed a new wave of work. In particular, an effort to confront semantic frameworks with empirical results has been made. In parallel, a number of mathematical advances have been made in modal logic, an area closely related to conditional logics. However, the techniques developed in modal logic with respect to proof theory and correspondence theory have not fully been applied to the conditional logics introduced in artificial intelligence and philosophy. The main objective of this seminar was to provide an opportunity for computer scientists, logicians, psychologists, linguists and philosophers working on that topic to meet and reinforce their ties over several days in the Dagstuhl castle.

We focused on three specific issues which were discussed and worked out in three different working groups. First, we investigated how people's intuitions about "counterpossibles" can be understood empirically and classified thanks to the theoretical accounts of conditional logics. Second, we reconsidered the various semantics of the basic system P and wondered to which extent pragmatics plays a role in the relevance relation between the antecedant and the consequent of a conditional. Third, we strove to apply the recent advances in proof theory and correspondence theory to conditional logics. These three topics correspond respectively to the working groups "Investigating people's intuitions about counterpossibles" (Section 4.1 of the full report), "The semantics of conditionals" (Section 4.2 of the full report) and "Correspondence theory and proof theory for conditional logics" (Section 4.3 of the full report).

These working group discussions were preceded by 13 short talks and 3 tutorials: "Semantics of Conditionals" (by Graham Priest), "Proof Theory of Conditionals" (by Nicola Olivetti) and "The psychology of Indicative Conditionals" (by Karolina Krzyzanowska). These talks and tutorials are summarized in Section 3 of the full report.

6.4 New Horizons in Parameterized Complexity

Organizers: Fedor V. Fomin, Dániel Marx, Saket Saurabh, and Meirav Zehavi

Seminar No. 19041

Date: January 20–25, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.1.67

Participants: Akanksha Agrawal, Andreas Björklund, Sergio Cabello, Parinya Chalermsook, Yijia Chen, Radu Curticapean, Holger Dell, Friedrich Eisenbrand, Henning Fernau, Fedor V. Fomin, Archontia C. Giannopoulou, Petr A. Golovach, Gregory Z. Gutin, Danny Hermelin, Yoichi Iwata, Bart Jansen, Mark R. Jerrum, Eun Jung Kim, Christian Komusiewicz, Martin Koutecký, Stefan Kratsch, Ariel Kulik, Euiwoong Lee, Bingkai Lin, Andrea Lincoln, Daniel Lokshtanov, Pasin Manurangsi, Dániel Marx, Matthias Mnich, Jesper Nederlof, Rolf Niedermeier, Shmuel Onn, Fahad Panolan, Marcin Pilipczuk, Michal Pilipczuk, M. S. Ramanujan, Peter Rossmanith, Ignasi Sau Valls, Saket Saurabh, Hadas Shachnai, Roohani Sharma, Stefan Szeider, Dimitrios M. Thilikos, Magnus Wahlström, Robert Weismantel, Meirav Zehavi

In 2019 the parameterized complexity (PC) community is celebrating two round dates: 30 years since the appearance of the paper of Abrahamson, Ellis, Fellows, and Mata in FOCS 1989, which can be considered as the starting point of PC, and 20 years since the appearance of the influential book of Downey and Fellows "*Parameterized Complexity*".

In these three decades, there has been tremendous progress in developing the area. The central vision of Parameterized Complexity through all these years has been to provide the algorithmic and complexity-theoretic toolkit for studying multivariate algorithmics in different disciplines and subfields of Computer Science. To achieve this vision, several algorithmic and complexity theoretic tools such as polynomial time preprocessing, aka kernelization, color-coding, graph-decompositions, parameterized integer programming, iterative compression, or lower bounds methods based on assumptions stronger than P=NP have been developed. These tools are universal as they did not only help in the development of the core of Parameterized Complexity, but also led to its success in other subfields of Computer Science such as Approximation Algorithms, Computational Social Choice, Computational Geometry, problems solvable in P (polynomial time) to name a few.

All cross-discipline developments result in flow of ideas and methods in both directions. In the last few years, we have witnessed several exciting developments of new parameterized techniques and tools in the following subfields of Computer Science and Optimization: Mathematical Programming, Computational Linear Algebra, Computational Counting, Derandomization, and Approximation Algorithms. A natural question is whether these domain-centric methods and tools are universal. That is, can they permeate boundaries of subfields and be employed wherever Parameterized Complexity approach can be used? The main objective of the seminar was to initiate the discussion on which of the recent domain-specific algorithms and complexity advances can become useful in other domains.

The seminar collected 46 participants from 18 countries. The participants presented their recent results in 26 invited and contributed talks. Open problems were discussed in open problem and discussion sessions.

6.5 Practical Yet Composably Secure Cryptographic Protocols

Organizers: Jan Camenisch, Ralf Küsters, Anna Lysyanskaya, and Alessandra Scafuro

Seminar No. 19042

Date: January 20–25, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.1.88

Participants: Jan Camenisch, Ran Canetti, Celine Chevalier, Ran Cohen, Manu Drijvers, Marc Fischlin, Dov Gordon, Jens Groth, Timo Hanke, Dennis Hofheinz, Markulf Kohlweiss, Stephan Krenn, Ralf Küsters, Anna Lysyanskaya, Mary Maller, Ueli Maurer, Arpita Patra, Antigoni Polychroniadou, Daniel Rausch, Alessandra Scafuro, Daniel Slamanig, Björn Tackmann, Muthuramakrishnan Venkitasubramaniam, Ivan Visconti, Sophia Yakoubov, Vassilis Zikas

We began by having survey talks on four research threads that had laid foundations of such models. Specifically, Ran Canetti presented his Universal Composability model, Dennis Hofheinz presented his work on the GNUC model, Ralf Küsters presented his IITM/iUC model, and Ueli Maurer presented the model of Constructive Cryptography.

Following these tutorials, we had several talks on how specific security goals and protocols are modeled and proved secure. Björn Tackmann presented a way to model a zero-knowledge proof protocol that made statements about knowledge of certain inputs to ideal functionalities. Manu Drijvers presented a way to model the global random oracle that can be used by participants in different protocols in a composable way.

Once the details of the specific models and how to use them were fresh in everyone's minds, we split up into working groups. In order to do this, we first had a discussion on what problems we believed were worth tackling; we proposed many problems, and then agreed to discuss a subset of them.

The topics explored by the working groups are discussed in detail below, in the "results" section of this report. The following additional topics were proposed for discussion (but were not discussed):

- Model asynchrony and time
- Anonymous communication
- Global random oracles in CC
- Secure Message Transfer in various model
- Concrete security in UC/IITM
- Finalise $\mathcal{F}_{\text{sig}}$ (with reasons why certain choices are better than others)

Additionally, we had several talks on recent and ongoing research projects. Marc Fischlin on composition of key agreement; Markulf Kohlweiss on structuring game-based proofs; Ran Cohen on probabilistic termination in cryptographic protocols; Antigoni Polychandrou presented two-round two-party computation; Vassilis Zikas modeling the public ledger functionality; Ran Canetti talking about using the EasyCrypt software to aid in cryptographic proofs and verification.

The following is a summary of the workshop results:

1. The relationship between the UC and IITM model was intensively discussed, concluding that the models are very close and that it is possible to unify the two models. The unification also seamlessly includes JUC, GUC, and SUC.
2. The working group on SNARKs (recursive composition of succinct proofs) achieved initial modeling success and crystallization of what's actually challenging.
3. The working group on modeling $\mathcal{F}_{\text{vrf}}$ and constricting it from $\mathcal{F}_{\text{sig}}$, $\mathcal{F}_{\text{ro}}$ figured out what the stumbling blocks were and what was fundamental.
4. The working group on $\mathcal{F}_{\text{NIZK}}$ and proofs about signatures in Constructive Crypto started to model typical UC functionality in the Constructive Crypto framework and then inspected how they could be composed.
5. The working group on building threshold primitives from single primitive (e.g. threshold signatures from signatures, threshold encryption from encryption etc) came up with a candidate for a "thresholdizer" functionality, and found some subtleties in defining threshold behavior in the ideal world. The also found a candidate construction to test the validity of the definition.
6. The working group on setup assumptions analyzed the assumptions used for constructing composable protocols in terms of practicality and security provided.
7. The working group on delegating secret keys – discovered a simple interface that can be added to $\mathcal{F}_{\text{sig}}$ to make it possible to delegate from one user to another well-defined user. Next steps are to investigate if it generalizes to other functionalities and to delegation that's based on knowledge transfer rather than explicit authorization of identity.

6.6 Data Structures for the Cloud and External Memory Data

Organizers: Gerth Stølting Brodal, Ulrich Carsten Meyer, Markus E. Nebel, and Robert Sedgewick

Seminar No. 19051

Date: January 27– February 1, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.1.104

Participants: Peyman Afshani, Hannah Bast, Michael A. Bender, Ioana Bercea, Timo Bingmann, Gerth Stølting Brodal, Alexander Conway, Martin Dietzfelbinger, Guy Even, Tomer Even, Rolf Fagerberg, Martin Farach-Colton, Johannes Fischer, Mordecai Golin, Herman J. Haverkort, John Iacono, Riko Jacob, Rob Johnson, Tsvi Kopelowitz, Moshe Lewenstein, Jérémie Lumbroso, Conrado Martinez, Kurt Mehlhorn, Ulrich Carsten Meyer, Friedhelm Meyer auf der Heide, Ian Munro, Robert Muth, Markus E. Nebel, John D. Owens, Manuel Penschuck, Seth Pettie, Sharath Raghvendra, Rajeev Raman, Alejandro Salinger, Peter Sanders, Robert Sedgewick, Siddhartha Sen, Francesco Silvestri, Robert Endre Tarjan, Sharma V. Thankachan, Sebastian Wild, Da Yan, Norbert Zeh

About the Seminar

Data structures provide ways of storing and manipulating data and information that are appropriate for the computational model at hand. Every such model relies on assumptions that we have to keep questioning. The aim of this seminar was to exchange ideas for new algorithms and data structures, and to discuss our models of computations in light of recent technological advances. This Dagstuhl seminar was the 13th in a series of loosely related Dagstuhl seminars on data structures.

Topics

The presentations covered both advances in classic fields, as well as new problems and insights for recent trends in computing. In particular, Johnson (Section 3.12 of the full report) and Muth (Section 3.17 of the full report) reported on models and research opportunities in the cloud and external memory motivated by practical demands.

A number of talks highlighted technical challenges in storing and processing large datasets: Bast (Section 3.2 of the full report) demonstrated the knowledge database QLever and discussed algorithmic aspects. Distributed frameworks were presented by Bingmann (Section 3.4 of the full report) reporting on the progress of Thrill while focusing on parallel external sorting and by Yan (Section 3.32 of the full report) who introduced G-thinker. Farach-Colton (Section 3.7 of the full report) analyzed the slow-down of various filesystems caused by updates over time. Owens (Section 3.19 of the full report) discuses intricacies of GPUs and presented efficient and practical data structures for this hardware.

In order to mitigate the impact of huge datasets, streaming and online algorithms were considered. Martinez (Section 3.15 of the full report) discussed Affirmative Sampling which takes uniform samples of a stream and adapts the sample size to the stream's diversity. Sedgewick (Section 3.26 of the full report) revisited the cardinality estimation problem and proposed the HyperBitBit algorithm. A matching of requests to resources in an online setting was covered by Raghvendra (Section 3.22 of the full report). Similarly, Mehlhorn (Section 3.16 of the full report) presented a solution to assigning indivisible resources approximately optimizing the social welfare.

Nebel (Section 3.18 of the full report) and Wild (Section 3.31 of the full report) proposed and analyzed tree-based data structures. Additionally, various aspects on more general graph processing were covered ranging from their enumeration (Lumbroso, Section 3.14 of the full report) and random sampling (Penschuck, Section 3.20 of the full report), over representations for k-connectivity (Pettie, Section 3.21 of the full report) to the detection of substructures (Silvestri, Section 3.28 of the full report and Tarjan, Section 3.29 of the full report).

Regarding the complexity of graph algorithms, Fagerberg (Section 3.23 of the full report) presented new lower bounds on the reorganisation cost of B-trees, while Thankachan (Section 3.30 of the full report) gave hardness results on the recognizability of Wheeler graphs. Kopelowitz (Section 3.13 of the full report) considered the complexity of data structures for the set-disjointness problem. Emphasizing cloud-related security concerns, Jacob (Section 3.11 of the full report) showed that a range of simple data structures have to incur an $\Omega(\log n)$ overhead if one wants to prevent information leakage via their access patterns.

Problems involving large text corpora were considered by Fischer (Section 3.8 of the full report) presenting an external memory bi-directional compression scheme, by Golin (Section 3.9 of the full report) discussing AIFV codes, and by Salinger (Section 3.24 of the full report) analyzing persistent full-text indices for versioned documents.

Data structures using hashing were examined by Conway (Section 3.5 of the full report), Dietzfelbinger (Section 3.6 of the full report), Even and Sanders (Section 3.25 of the full report). Bender (Section 3.3 of the full report) discussed variants of Bloom filters which adapt based on past queries.

Afshani (Section 3.1 of the full report) presented Fragile Complexity, a novel model of computation with an element-centric cost function, and gave bounds for various classical problems. Iacono (Section 3.10 of the full report) proposed to model locality-of-reference more explicitly and compared his proposal to the external memory and cache-oblivious model. Sen (Section 3.27 of the full report) proposed the novel paradigm HAIbrid augmenting classic data structures with artificial intelligence.

Final Thoughts

The organizers would like to thank the Dagstuhl team for their continuous support; the welcoming atmosphere made the seminar both highly productive and enjoyable. They also thank all participants for their contributions to this seminar.

6.7 Computational Methods for Melody and Voice Processing in Music Recordings

Organizers: Meinard Müller, Emilia Gómez, and Yi-Hsuan Yang

Seminar No. 19052

Date: January 27–February 1, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.1.125

Participants: Rachel Bittner, Estefanía Cano Cerón, Michèle Castellengo, Pritish Chandna, Helena Cuesta, Johanna Devaney, Simon Dixon, Zhiyao Duan, Emilia Gómez, Masataka Goto, Frank Kurth, Cynthia Liem, Antoine Liutkus, Meinard Müller, Tomoyasu Nakano, Juhan Nam, Ryo Nishikimi, Geoffroy Peeters, Polina Proutskova, Preeti Rao, Sebastian Rosenzweig, Justin Salamon, Frank Scherbaum, Sebastian J. Schlecht, Li Su, Tomoki Toda, Julián Urbano, Anja Volk, Ye Wang, Christof Weiß, Yi-Hsuan Yang, Frank Zalkow

In this executive summary, we give an overview of computational melody and voice processing and summarize the main topics covered in this seminar. We then describe the background of the seminar's participants, the various activities, and the overall organization. Finally, we reflect on the most important aspects of this seminar and conclude with future implications and acknowledgments.

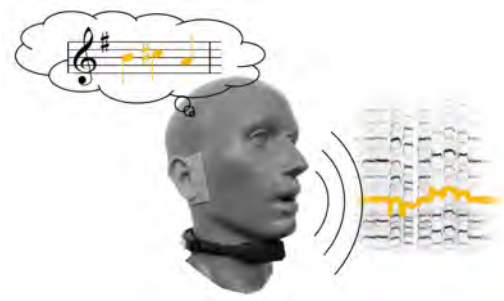

Overview

When asked to describe a specific piece of music, we are often able to sing or hum the main melody. In general terms, a *melody* may be defined as a linear succession of musical tones expressing a particular musical idea. Because of the special arrangement of tones, a melody is perceived as a coherent entity, which gets stuck in a listener's head as the most memorable element of a song. As the original Greek term *melōidía* (meaning "singing" or "chanting") implies, a melody is often performed by a human voice. Of course, a melody may also be played by other instruments such as a violin in a concerto or a saxophone in a jazz piece. Often, the melody constitutes the leading element in a composition, appearing in the foreground, while the accompaniment is in the background. Sometimes melody and accompaniment may even be played on a single instrument such as a guitar or a piano. Depending on the context and research discipline (e. g., music theory, cognition or engineering), one can find different descriptions of what may be meant by a melody. Most people would agree that the melody typically stands out in one way or another. For example, the melody often comprises the higher notes in a musical composition, while the accompaniment consists of the lower notes. Or the melody is played by some instrument with a characteristic timbre. In some performances, the notes of a melody may feature easily discernible time–frequency patterns such as vibrato, tremolo, or glissando. In particular, when considering performed music given in the form of audio signals, the detection, extraction, separation, and analysis of melodic voices becomes a challenging research area with many yet unsolved problems. In the following, we discuss some MIR tasks related to melody processing, indicating their relevance for fundamental research, commercial applications, and society.

The problem of detecting and separating melodic voices in music recordings is closely related to a research area commonly referred to as *source separation*. In general, audio signals are complex mixtures of different sound sources. The sound sources can be several people talking simultaneously in a room, different instruments playing together, or a speaker talking in the foreground with music being played in the background. The general goal of source separation is to decompose a complex sound mixture into its constituent components. Source separation methods often rely on specific assumptions such as the availability of multiple channels, where several microphones have been used to record the acoustic scene from different directions. Furthermore, the source signals to be identified are assumed to be independent in a statistical sense. In music, however, such assumptions are not applicable in many cases. For example,

musical sound sources may outnumber the available information channels, such as a string quartet recorded in two-channel stereo. Also, sound sources in music are typically highly correlated in time and frequency. Instruments follow the same rhythmic patterns and play notes which are harmonically related. This makes the separation of musical voices from a polyphonic sound mixture an extremely difficult and generally intractable problem.

When decomposing a music signal, one strategy is to exploit music-specific properties and additional musical knowledge. In music, a source might correspond to a melody, a bass line, a drum track, or a general instrumental voice. The separation of the melodic voice, for example, may be simplified by exploiting the fact that the melody is often the leading voice, characterized by its dominant dynamics and by its temporal continuity. The track of a bass guitar may be extracted by explicitly looking at the lower part of the frequency spectrum. A human singing voice can often be distinguished from other musical sources due to characteristic time–frequency patterns such as vibrato. Besides such acoustic cues, score-informed source separation strategies make use of the availability of score representations to support the separation process. The score provides valuable information in two respects. On the one hand, pitch and timing of note events provide rough guidance within the separation process. On the other hand, the score offers a natural way to specify the target sources to be separated.

In this seminar, we discussed source separation techniques that are particularly suited for melodic voices. To get a better understanding of the problem, we approached source separation from different directions including model-based approaches that explicitly exploit acoustic and musical assumptions as well as data-driven machine learning approaches.

Given a music recording, melody extraction is often understood in the MIR field as the task of extracting a trajectory of frequency values that correspond to the pitch sequence of the dominant melodic voice. As said before, melody extraction and source separation are highly related: while melody extraction is much easier if the melodic source can be isolated first, the source separation process can be guided if the melodic pitch sequence is given a priori. However, both tasks have different goals and involve different challenges. The desired output of melody extraction is a trajectory of frequency values, which is often sufficient information for retrieval applications (e. g., query-by-humming or the search of a musical theme) and performance analysis. In contrast, for music editing and audio enhancement applications, source separation techniques are usually needed.

In the seminar, we addressed different problems that are related to melody extraction. For example, the melody is often performed by a solo instrument, which leads to a problem also known as *solo–accompaniment separation*. The estimation of the fundamental frequency of a quasi-periodic signal, termed *mono-pitch estimation*, is a long-studied problem with applications in speech processing. While mono-pitch estimation is now achievable with reasonably high accuracy, the problem of *multi-pitch estimation* with the objective of estimating the fundamental frequencies of concurrent periodic sounds remains very challenging. This particularly holds for music signals, where concurrent notes stand in close harmonic relation. For extreme cases such as complex orchestral music where one has a high level of polyphony, multi-pitch estimation becomes intractable with today's methods.

Melodic voices are often performed by singers, and the singing voice is of particular importance in music. Humans use singing to create an identity, express their emotions, tell stories, exercise creativity, and connect while singing together. Because of its social, cultural, and educational impact, singing plays a central role in many parts of our lives, it has a positive effect on our health, and it creates a link between people, disciplines, and domains (e. g., music and language). Many people are active in choirs, and vocal music makes up an important part of our cultural heritage. In particular in Asian countries, karaoke has become a major cultural force performed by people of all age groups. Singing robots, vocaloids, or synthesizers such as Hatsune Miku[43] have made their way into the mass market in Japan. Thanks to digitization and technologies, the world wide web has become an important tool for amateur and professional singers to discover and study music, share their performances, get feedback, and engage with their audiences. An ever-increasing amount of music-related information is available to singers and singing enthusiasts, such as music scores[44] as well as audio and video recordings.[45] Finally, music archives contain an increasing number of digitized audio collections of historic value from all around the world such as Flamenco music, Indian art music, Georgian vocal music, or Bejing Opera performances.

Due to its importance, we placed in our seminar a special emphasis on music technologies related to singing. This involves different research areas including singing analysis, description, and modeling (timbre, intonation, expression), singing voice synthesis and transformation, voice isolation/separation, and singing performance rating. Such research areas require a deep understanding of the way people produce and perceive vocal sounds. In our seminar, we discussed such issues with researchers having a background in singing acoustics and music performance.

Over the last years, as is also the case for other multimedia domains, many advances in music and audio processing have benefited from new developments in machine learning.

In particular, deep neural networks (DNNs) have found their way into MIR and are applied with increasing success to various MIR tasks including pitch estimation, melody extraction, sound source separation, and singing voice synthesis. The complex spectro-temporal patterns and relations found in music signals make this domain a challenging testbed for such new machine learning techniques. Music is different from many other types of multimedia. In a static image, for example, objects may occlude one another with the result that only certain parts are visible. In music, however, concurrent musical events may superimpose or blend each other in a more complicated way. Furthermore, as opposed to static images, music depends on time. Music is organized in a hierarchical way ranging from notes, bars, and motifs, to entire sections. As a result, one requires models that capture both short-term and long-term dependencies in music.

In the seminar, we looked at the new research challenges that arise when designing music-oriented DNN architectures. Furthermore, considering the time-consuming and labor-intensive process of collecting human annotations of musical events and attributes (e. g., timbre, intonation, expression) in audio recordings, we addressed the issue of gathering large-scale annotated datasets that are needed for DNN-based approaches.

[43] https://en.wikipedia.org/wiki/Hatsune_Miku

[44] For example, the Choral Public Domain Library currently hosts free scores of at least 24963 choral and vocal works by at least 2820 composers, see http://www.cpdl.org/

[45] See, for example, the material hosted at platforms such as YouTube or SoundCloud

Participants and Group Composition

In our seminar, we had 32 participants, who came from various locations around the world including North America (4 participants from the U.S.), Asia (4 participants from Japan, 2 from Taiwan, 2 from Singapore, 1 from Korea, 1 from India), and Europe (18 participants from France, Germany, Netherlands, Spain, United Kingdom). More than half of the participants came to Dagstuhl for the first time and expressed enthusiasm about the open and retreat-like atmosphere. Besides its international character, the seminar was also highly interdisciplinary. While most of the participating researchers are working in the field of music information retrieval, we also had participants with a background in musicology, acoustics, machine learning, signal processing, and other fields. By having experts working in technical as well as in non-technical disciplines, our seminar stimulated cross-disciplinary discussions, while highlighting opportunities for new collaborations among our attendees. Most of the participants had a strong musical background, some of them even having a dual career in an engineering discipline and music. This led to numerous social activities including singing and playing music together. In addition to geographical locations and research disciplines, we tried to foster variety in terms of seniority levels and presence of female researchers. In our seminar, 10 of the 32 participants were female, including three key researchers (Anja Volk, Emilia Gómez, and Johanna Devaney) from the "Women in Music Information Retrieval" (WiMIR)[46] initiative.

In conclusion, by gathering internationally renowned scientists as well as younger promising researchers from different research areas, our seminar allowed us to gain a better understanding of the problems that arise when dealing with a highly interdisciplinary topic such as melody and voice processing–problems that cannot be addressed by simply using established research in signal processing or machine learning.

Overall Organization and Schedule

Dagstuhl seminars are known for having a high degree of flexibility and interactivity, which allows participants to discuss ideas and to raise questions rather than to present research results. Following this tradition, we fixed the schedule during the seminar asking for spontaneous contributions with future-oriented content, thus avoiding a conference-like atmosphere, where the focus tends to be on past research achievements. After the organizers gave an overview of the Dagstuhl concept and the seminar's overall topic, we started the first day with self-introductions, where all participants introduced themselves and expressed their expectations and wishes for the seminar. We then continued with a small number of short (15 to 20 minutes) stimulus talks, where specific participants were asked to address some critical questions on melody and voice processing in a nontechnical fashion. Each of these talks seamlessly moved towards an open discussion among all participants, where the respective presenters took over the role of a moderator. These discussions were well received and often lasted for more than half an hour. The first day closed with a brainstorming session on central topics covering the participants' interests while shaping the overall schedule and format for the next day. On the subsequent days, we continued having stimulus tasks interleaved with extensive discussions. Furthermore, we split into smaller groups, each group discussing a more specific topic in greater depth. The results and conclusions of these parallel group sessions, which lasted between 60 to 90 minutes, were then presented and discussed with the plenum. This mixture of presentation elements gave all participants the opportunity for presenting their ideas while avoiding a monotonous conference-like presentation format. On the last day, the seminar concluded with a session we called "self-outroductions" where each participant presented his or her personal view on the seminar's results.

Additionally to the regular scientific program, we had several additional activities. First, we had a demo session on Thursday evening, where participants presented user interfaces, available datasets, and audio examples of synthesized singing voices. One particular highlight was the incorporation of singing practice in the seminar. In particular, we carried out a recording session on Wednesday afternoon, where we recorded solo and polyphonic singing performed by Dagstuhl participants. The goal of this recording session was to contribute to existing open datasets in the area of music processing. The singers were recorded with different microphone types such as throat and headset microphones to obtain clean recordings of the individual voices. All participants agreed that the recorded dataset should be made publicly available for research purposes. As preparation for these recordings, we assembled a choir consisting of ten to twelve amateur singers (all Dagstuhl participants) covering different voice sections (soprano, alto, tenor, bass). In the lunch breaks and the evening hours, the group met for regular rehearsals to practice different four-part choral pieces. These musical activities throughout the entire week not only supported the theoretical aspects of the seminar but also had a very positive influence on the group dynamics. Besides the recordings, we also had a concert on Thursday evening, where various participant-based ensembles performed a variety of music including classical music and folk songs.

Conclusions and Acknowledgment

Having a Dagstuhl seminar, we gathered researchers from different fields including information retrieval, signal processing, musicology, and acoustics. This allowed us to approach the problem of melody and voice processing by looking at a broad spectrum of data analysis techniques (including signal processing, machine learning, probabilistic models, user studies), by considering different domains (including text, symbolic, image, audio representations), and by drawing inspiration from the creative perspectives of the agents (composer, performer, listener) involved. As a key result of this seminar, we achieved some substantial progress towards understanding, modeling, representing, and extracting melody- and voice-related information using computational means.

The Dagstuhl seminar gave us the opportunity for having interdisciplinary discussions in an inspiring and retreat-like atmosphere. The generation of novel, technically oriented scientific contributions was not the main focus of the seminar. Naturally, many of the contributions and discussions were on a conceptual level, laying the foundations for future projects and collaborations. Thus, the main impact of the seminar is likely to take place in the medium and long term. Some more immediate results, such as plans to share research data and software, also arose from the discussions. In particular, we plan to make the dataset recorded during the Dagstuhl seminar available to the research community.

[46] https://wimir.wordpress.com/

As further measurable outputs from the seminar, we expect to see several joint papers and applications for funding.

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

In conclusion, our expectations for the seminar were not only met but exceeded, in particular concerning networking and community building. We want to express our gratitude to the Dagstuhl board for giving us the opportunity to organize this seminar, the Dagstuhl office for their exceptional support in the organization process, and the entire Dagstuhl staff for their excellent service during the seminar. In particular, we want to thank Susanne Bach-Bernhard, Annette Beyer, Michael Gerke, and Michael Wagner for their assistance during the preparation and organization of the seminar.

Fig. 6.1
"#Dagstuhl Seminar = sticky notes + coffee + multilayer networks + coffee + colleagues from all disciplines + coffee." Twitter post by 19061 Dagstuhl Seminar participant Martin Grandjean.
https://twitter.com/GrandjeanMartin/status/1093543791691288578. Photo courtesy of Martin Grandjean.

6.8 Visual Analytics of Multilayer Networks Across Disciplines

Organizers: Mikko Kivelä, Fintan McGee, Guy Melançon, Nathalie Henry Riche, and Tatiana von Landesberger

Seminar No. 19061

Date: February 3–8, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.2.1

Participants: Jan Aerts, Daniel Archambault, Ariful Azad, Kathrin Ballweg, Fabian Beck, Remco Chang, Alberto Cottica, Marten Düring, Jean-Daniel Fekete, Mohammad Ghoniem, Martin Grandjean, Jessie Kennedy, Andreas Kerren, Mikko Kivelä, Søren Knudsen, Stephen G. Kobourov, Sylvain Legay, Matteo Magnani, Maria Malek, Mark McCann, Fintan McGee, Guy Melançon, Anders Kristian Munk, Bruno Pinaud, Margit Pohl, Helen C. Purchase, Benjamin Renoust, Arnaud Sallaberry, Christian Tominski, Paola Valdivia, Jason Vallet, Tatiana von Landesberger, Michael Wybrow, Björn Zimmer, Blaž Škrlj

Introduction

The topic of multilayer networks has recently emerged from the field of complex systems, however many of the of the fundamental concepts and ideas have existed for some time, in fields such as sociology, and often under different nomenclature, such as multimodal, heterogeneous or multiplex networks. The multilayer network framework of Kivelä *et al* [1] has collected many of these concepts and different labels, along with example data sets, allowing us to recognize the multi-disciplinary importance of multilayer networks as a topic. Despite the importance of this topic, it is only recently that the visualization community is beginning to consider approaches for the visual analytics of multilayer networks. This seminar was the first to bring together practitioners from multiple domains to discuss the visual analytics of multilayer networks. These fields included data visualization, complex systems, digital humanities, biological sciences, health informatics, and sociology. The primary goal of this seminar was to bring together these thinkers and practitioners from different disciplines to drive forward new advances on the topic. The seminar was designed to foster discussions between researchers and designers of visual analytics tools, those who define the underlying theory, and the the end-users of these tools. To push research further and produce significant impact in industry and general public practices, the research community needs to establish a deeper collaboration between data scientists and researchers from applications domains (e.g. biologists, social scientists, business analysts, journalists, physicists), who collect and analyze the data; and researchers in maths, physics and computer science who push the state-of-the-art, producing visualization and analysis models, algorithms and tools. This deeper collaboration starts with building an understanding of the needs and tasks of network analysts. This seminar was an important first step, leveraging cross domain synergies with the goal of identifying the shared underlying problems and helping to solve them. The domain experts presented their domain problems early on in the seminar, and then interacted with two different sets of visualization experts in two separate breakout sessions. The motivation for this was to expose the visualization experts to many different domain problems and to expose the domain experts to multiple approaches to their problems. Our goal was to not only to advance research in the field of visualization, but also to provide techniques to help the domain experts to advance research in their own field. Interdisciplinary intersection was a key part of the methodology of our seminar.

Seminar Topics

The seminar featured talks and working groups that discussed topics on visualization, analysis, theory and applications of multilayer networks (see Sections 3 and 4). The application domain focus was maintained throughout the seminar. Experts from application domains gave talks in the first day and a half highlighting the problems they encountered. Then there were two breakout sessions where each experts was assigned a different group of visualization experts, allowing the domain experts to brainstorm solutions to their problems with different sets of visualization experts.

Talks. The talks brought the interdisciplinary participants initial information on a) current application problems dealt with in the area of multilayer networks and b) current visualization, analysis and systems solutions.

The purpose of talks by application experts was to make sure that the potential solutions provided by the interactive visualizations and analytics fully meet the requirements of those who actually use them, i.e. the system biologists, social network analysts, historians, etc. Therefore, the talks provided understanding of the data and problems/tasks/goals when analyzing multilayer graphs by the domain experts. The talks covered

application areas of social networks by A. Cottica and by M. Magnani, information circulation in an international organization by M. Grandjean, digital humanities by M. Düring , multi-omics data by S. Legay, population health by M. McCann, digital ethnography by A. Munk(see Section 3 of the full report). These talks allowed the visualizations and complex systems theory experts to gain some insight into the domain experts problems. As we also wanted the domain experts to be exposed to multiple approaches to their problems, we had two breakout sessions after all of the domain experts presented their personal topics. In these breakout sessions, each domain experts was assigned a small group of visualization researchers to further brainstorm, mapping visualization problems to domain problems. Different researchers were assigned to each domain expert for each session. This exposed the domain experts to multiple visualization approaches, and allowed for synergies between application domain problems to be identified by the visualization researchers. At the end of each breakout session, each groups gave a short report back to other participants, allowing for further discussion and cross fertilization of ideas. This approach ensured that that both the domain experts and the visualization experts had a wide range of ideas to explore as part of the working groups in the later half of the seminar.

The purpose of talks by visualization, analysis and systems experts was to present currently available solutions to multilayer network visualization and analysis (see also Section 3 of the full report). These talks were dispersed throughout out the week. The talk topics were: an introduction to multilayer networks by F. McGee, a complex systems perspective on the concept of multilayer network by M. Kivelä, survey of multilayer visualizations by G. Melançon, Py3plex library for visualization by B. Škrlj, interaction with multilayer network visualization by B. Renoust. This allowed application experts to get to know the advantages and limitations of existing solutions. The talk schedule was flexible, for example, due to a high level of interest form all attendees M. Kivelä gave a second question and answers session to his talk the following day.

Working Groups. At the midpoint of the week we defined the working groups. The breakout session stimulated a large amount of discussion and ideas across participants of all disciplines. Following on the breakout sessions discussions, all seminar participants wrote down topics and ideas that were that were of interest to them of pieces of paper, which were affixed to a board. Similar topics were re-positioned closer together on the board, until all participants reached a consensus of five topic areas for discussion within working groups. The resulting working groups were as follows:

- *Unifying Terminology and visual analytic approaches*: One open problem of multilayer network analysis and visualization is the inconsistent terminology across disciplines. There are many different names given to networks with such characteristics, outlining the current lack of consistent definitions between disciplines, such as heterogeneous, multi-faceted, multi-modal, or multi-relational networks, amongst others (see [1]) and in the vast majority of cases it is possible to model them as multilayer networks. The discussion group assessed various types of networks from visualization, application and systems perspective. It discussed possible unification of these perspectives in one visual analytics framework and identified open challenges (see Section 4.3 of the full report).
- *Analytics, Communities Comparison and attributes*: Visual analysis of multilayer networks is also concerned with the exhibition of salient properties and patterns in data. Salience in networks is often captured through metrics (networks statistics) while patterns most often correspond to particular subsets of entities (nodes and edges). Layers bring additional complexity to the computation of these metrics and patterns, as metrics and patterns may need to be computed across several layers. The visualization of the computed metrics and patterns needs to consider also these layers, thus, posing challenges to the data presentation. This working group analyzed the current network metrics and proposed novel metrics specifically for multilayer networks (see Section 4.4 of the full report).
- *Interaction (and Layer Creation)*: (see Section 4.2 of the full report) This topic concentrated on interactive creation of layers in networks. While the input multilayer network may have predefined layers, in many use cases, the layers need to be adapted to the analytical task during network exploration. This working group has gathered requirements for interaction with layers, surveyed current solutions and their limitations. They have proposed novel approaches that will be pursued after the seminar.
- *Visual Encodings* The complex relationships between complex structures mean that traditional interactive visualizations need to be enhanced. Researchers from the various domains can exchange their ideas and thus start novel avenues in interactive visualization. The discussion of this working group focused on the visualization design – encodings. The group identified main requirements for visualization: aggregations, interactive layer editing, overview of all layers, details of an individual layer and exploration paths – top-down versus bottom up (see Section 4.5 of the full report). These requirements are used to derive a design space of possible visualization approaches in future.
- *Human Factors and Multilayer Networks* This topic focused on the user's point of view in the design of multilayer network visualization. This is a challenge as the complexity of multilayer networks results in a significant amount of cognitive load on the users. The group collected results from related work that can be used as guidelines for designing multilayer visualizations. It also identified gaps in literature for future research (see Section 4.1 of the full report).

Seminar Outcomes

During the seminar, a number of sub-topics were identified that require further research: *A unifying visualization framework*, *Novel Visual Encodings*, *Analytics* and *Attributes*, *Interaction*, *Evaluation*, *Use Cases* and *Human Factors*.

- **A unifying visualization framework for multilayer networks:** Currently, multilayer networks are referred to across communities using various names and concepts. A novel unified conceptual framework for multilayer network is needed that would be used for visualization, interaction and analytics purposes. It should extend the underlying mathematical framework [1] to meet the needs of the data and tasks associated with the various use cases, as well as existing visualization and interaction concepts.
- **Novel visual encodings:** The existing visualization techniques have limited scope for the broad range of data and tasks in the applications of multilayer networks. Therefore, novel visual encodings need to be researched that to enable data exploration across layers.
- **Interaction:** Visual exploration and analysis of multilayer networks requires novel interaction techniques that would allow to browse across layers and also to create new layers during the exploration process.

- **Interdisciplinarity:** The wide range of application domain problems sets novel problems that may be best addressed by new visualization approaches. The development of novel solutions for visual analysis of multilayer networks requires joined forces of application, visualization and analysis experts.
- **Multiple layers and attributes:** The complexity of multilayer networks often includes an additional dimension: The multivariate nature of node and edge attributes. This information needs to be encoded in the visualization and supported in analytical functions. This raises novel challenges.
- **Network Analytics:** Visual network analysis also covers the understanding the analytical relationship between layers (with respect to structure and/or attributes) and the layer comparison. The limitations of current analytical approaches and network metrics raises many interesting challenges and opportunities for developing new metrics for the multilayer use case.
- **Evaluation & Human Factors:** The human perspective on the complexity of the network structure and its visualization needs to be assessed. It covers a) the perceptual and cognitive aspects when interactively exploring the networks and b) a thorough empirical evaluation of the analytical paths and insights. The existing methodologies for such research should be adapted for the multilayer network case.

These topics will be discussed in the follow-up VIS 2019 Workshop "Challenges in Multilayer Network Visualization and Analysis". The workshop is co-organized by Dagstuhl Seminar organizers and participants: Fintan McGee, Tatiana von Landesberger, Daniel Archambault and Mohammad Ghoniem. The seminar will feature keynote, paper and poster sessions as well as discussion rounds on the above-mentioned topics.

References

1 Mikko Kivelä and Alex Arenas and Marc Barthelemy and James P. Gleeson and Yamir Moreno and Mason A. Porter (1998). *Multilayer networks*. Journal of Complex Networks.

6.9 Bringing CP, SAT and SMT together: Next Challenges in Constraint Solving

Organizers: Sébastien Bardin, Nikolaj S. Bjørner, and Cristian Cadar

Seminar No. 19062

Date: February 3–6, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.2.27

Participants: Sébastien Bardin, Armin Biere, Nikolaj S. Bjorner, François Bobot, Frank Busse, Cristian Cadar, Maria Christakis, Bruno Dutertre, Benjamin Farinier, Pierre Flener, Sean Heelan, Matti Järvisalo, Timotej Kapus, Laura Kovács, Laurent Michel, Yannick Moy, Robert Nieuwenhuis, Jakob Nordström, Marie Pelleau, Mauro Pezzè, Tanja Schindler, Christian Schulte, Laurent Simon, Mate Soos, Peter J. Stuckey, Cesare Tinelli, Andrei Voronkov

The scattered landscape of constraint solving. Constraint solving is at the heart of several key technologies, including program analysis, testing, formal methods, compilers, security analysis, optimization, and AI. During the last two decades, constraint solving has been highly successful and transformative: on the one hand, SAT/SMT solvers have seen a significant performance improvement with a concomitant impact on software engineering, formal methods and security; on the other hand, CP solvers have also seen a dramatic performance improvement, with deep impact in AI and optimization.

These successes bring new applications together and new challenges: some fundamental constraints still lack efficient reasoning (e.g., floating-point arithmetic); quantifiers are rarely taken into account; current approaches focus essentially on satisfiability and/or validity while some applications would benefit from queries such as optimization or model counting. While each of the SAT, SMT and CP communities has made progress on some of these problems, no approach is able to tackle them all. Moreover, while historically strongly connected, the SAT/SMT communities have had minimal interactions with the CP community over the recent years.

Goals. The aim of this seminar was to reunify the Constraint Solving landscape and identify the next big challenges together with promising approaches. The seminar brought together researchers from SAT, SMT and CP along with applications researchers in order to foster cross-fertilization of ideas, deepen interactions, identify the best ways to serve the application fields and in turn help improve the solvers for specific usages.

An overview of constraint solving.

- **CP.** *Constraint Programming* [1] focuses on finding a solution (satisfiability) or a best solution (optimization) to constraint problems seen as sets of atomic constraints over arbitrary domains. Traditionally, CP is interested in problems defined over finite-domain variables (typically: bounded integers), yet a lot of work has also been devoted to infinite domains such as real numbers. The basic scheme of CP approaches (in the finite setting) consists in exploring the search tree of all partial valuations of the problem until a solution is found, or all possible valuations have been explored. At each step, *propagation* allows to refine further the admissible values for yet-unlabeled variables and, once no more propagation is possible, *labeling* assigns a value to a yet-unlabeled variable (yielding a backtrack point) and then propagation takes place against this, etc. CP has been highly successful in AI-related domains such as planning or scheduling, and promising applications to program verification have emerged recently.

  Strong points: advance propagation techniques based on the key notion of arc-consistency; specific reasoning, especially for finite-domain theories (e.g. floats, bounded arithmetic, bitvectors); queries beyond satisfiability, e.g. optimization

- **SAT.** While the seminal DPLL procedure [3] follows mostly the procedure described above for CP but specialized to the Boolean case[47], the true miracle of SAT comes from its modern version [2], where conflict-driven learning allows significant driven-by-need pruning of the search space—making the technique equally good at finding solutions or proving there is none. Many more improvements have been explored over the years, with carefully tuned propagation, data structures and branching heuristics. DPLL-style SAT solvers are at the core of hardware design and verification tools, and they have shown unreasonable efficiency on very large industrial problems.

[47] Seeing CP as a generalization of SAT is also possible.

Strong points: conflict-driven clause learning methods; efficient search/propagate procedure, with optimized branching and look-ahead.

- **SMT.** *Satisfiability Modulo Theory* [4] extends SAT by considering the satisfiability problem over combinations of first-order theories, for examples formulas involving complex boolean structure plus uninterpreted functions, arrays and linear arithmetics. While first restricted to the unquantified case, the technique has been extended with partial support for quantifiers. The core of SMT techniques is the combination of efficient theory-dedicated conjunctive-only decision procedures (typically through the Nelson-Oppen combination framework) together with their lifting to the general (disjunctive) case thanks to the DPLL(T) framework, where a DPLL-style SAT solver works in interplay with theory solvers. SMT problems arise naturally in software analysis, where programs are built over combinations of basic data types. Hence, SMT solvers are naturally at the heart of most modern software verification technologies.

  Strong points: first-order decision procedures, including theories over infinite domains; elegant combinations of solvers; partial handling of quantifiers.

Research questions. The seminar allows to highlight several key challenges to current constraint solving techniques. They have been discussed during the meeting from different research perspectives.

- Hard-to-handle data types: several common data types and associated theories are still not managed in an efficient-enough way, typically finite-but-large domains such as modular arithmetic, bounded arithmetic with non-linear operations, floating-point arithmetic or bitvector constraints deeply mixing arithmetic and bit-level reasoning, sets with cardinality, strings with size, etc.
- Quantifiers: quantifiers can be added to SMT solvers but often at the price of losing model generation, while there is some support for finite quantification in SAT and CP but at the price of a significant drop in performance; yet, quantifiers are useful in practice (initial state, pre/post-conditions, summaries, etc.);
- Beyond satisfiability: while the first applications of constraint solving were concerned with finding solutions or proving validity / infeasibility, new applications bring new types of queries, such a optimization, soft constraints, solution counting, over-approximating sets of solutions, etc.
- New trade-offs between learning and propagation: while the SAT community seems to have reached a sweet spot on this question (with efforts put on a posteriori learning rather than on a priori propagation), the issue is not settled yet for SMT and CP, and may be theory and/or application dependent.

Potential synergies. We have also identified the following potential synergies between CP, SAT and SMT, and expect strong interactions around these points in a near future:

- CP researchers have advanced propagation techniques, domain-dedicated reasoning and (deep) constraint combination. SAT and SMT researchers can learn from that.
- SAT researchers have significantly advanced branching heuristics, look-ahead and conflict-clause learning methods. CP and SMT researchers can learn from that.
- SMT researchers have focused on theory solvers and well-defined solver combinations. How can we do "lightweight" theory integration in SAT/CP solvers that trade off generality for cheaper and focused implementation of theories aimed at very specific applications? SAT and CP researchers can take advantage of these points.
- How can we better serve the needs of applications researchers? Application researchers can tell solver designers about which of these features (and combinations thereof) they would like the most in a single solver.
- Finally, an important question is how do we leverage machine learning in these contexts. The experience of the SAT community may bring here some answers.

Outcome. The main goal of this Dagstuhl seminar was to bring together leading researchers in the different subfields of automated reasoning and constraint solving, foster greater communication between these communities and discuss new research directions.

The seminar had 28 participants from Australia, Austria, France, Germany, Finland, Italy, Spain, Sweden, Switzerland, United Kingdom and United States, from both academia, research laboratories and the industry. More importantly, the participants represented several different communities, with the topics of the talks and discussions reflecting these diverse interests in both solving technologies (CP, SAT, SMT), challenges (floating-point constraints, quantifiers, etc.) and application domains (testing, verification, security, compilation, commercialization, among others).

It was the first time such an *inclusive* meeting was held, bringing together leading researchers from SAT/SMT (typical interest: formal verification), CP (typical interest: optimization) and applications (typical interest: testing, verification, security). All participants agreed the event was fruitful, and we expect to see more collaborations between SAT/SMT and CP in a near future.

References

1 R. Dechter. *Constraint processing*. Elsevier Morgan Kaufmann, 2003.

2 L. Zhang, C. F. Madigan, M. H. Moskewicz, and S. Malik. Efficient conflict driven learning in a boolean satisfiability solver. In *International Conference on Computer-aided design*. IEEE Press, 2001.

3 M. Davis, G. Logemann, and D. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 1962.

4 C. W. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability modulo theories. In *Handbook of Satisfiability*. 2009.

6.10 Specification Formalisms for Modern Cyber-Physical Systems

Organizers: Jyotirmoy V. Deshmukh, Oded Maler, and Dejan Ničković

Seminar No. 19071

Date: February 10–15, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.2.48

Participants: Houssam Abbas, Nikos Aréchiga, Ezio Bartocci, Marcello M. Bersani, Paul Bogdan, Borzoo Bonakdarpour, Chih-Hong Cheng, Thao Dang, Jyotirmoy Deshmukh, Rayna Dimitrova, Alexandre Donzé, Katie Driggs-Campbell, Georgios Fainekos, Lu Feng, Thomas Ferrère, Bernd Finkbeiner, Dana Fisman, Felipe Gorostiaga, Radu Grosu, James Kapinski, Martin Leucker, Rupak Majumdar, Niveditha Manjunath, Stefan Mitsch, Laura Nenzi, Dejan Ničković, Jens Oehlerking, Ana Oliveira da Costa, Necmiye Ozay, Pavithra Prabhakar, Gustavo Quirós, Akshay Rajhans, David Šafránek, César Sánchez, Caleb Stanford, Hazem Torfah, Marcell Vazquez-Chanlatte

Modern Cyber-Physical Systems (CPS) represent the convergence of the fields of control theory, artificial intelligence, machine learning, and distributed communication/coordination. CPS applications range from small quad-rotor based aerial vehicles to commercial airplanes, from driverless autonomous vehicles to vehicle platoons, from nano-scale medical devices in closed-loop with a human to giga-scale industrial manufacturing systems. While several application domains can claim to be cyber-physical systems, a unique aspect of CPS is a strong focus on model-based development (MBD). The MBD paradigm allows analyzing the system virtually, examining its safety, performance, stability, security, privacy, or resilience. At a certain level of abstraction, a model of a CPS application can be roughly divided into three parts: (1) the plant model representing an encapsulation of the physical components in the system, (2) the controller model representing the software used to regulate the plant, and (3) an environment model representing exogenous disturbances to the plant.

Given plant, controller and environment models of a system, in order to perform any of the aforementioned analyses, a crucial step is to articulate the goal of the analysis as a formal specification for the system. The analysis problem can then check whether the system implementation is a refinement of its specification. However, the state-of-the-art in industrial settings is that formal specifications are rarely found. Specifications exist in the form of mental models of correctness formed by engineers through their design insights and experience, or visual depictions in the form of simulation plots, and occasionally as legacy scripts and monitors. None of these are formal, machine-checkable unambiguous specifications. In the industry, engineers often use the term requirements instead of specifications. Typical industrial requirements do not arise from principled software engineering approaches to develop CPS software, but rather are summaries of discussions between developers and their customers. While the state-of-the-art for requirements/specifications in industrial settings is far from ideal, in academic settings, there is a problem of having a wide choice between a number of specification formalisms, primarily being developed by the formal methods community. On the other hand, application-specific academic domains such as robotics, biological systems, and medical devices may not always articulate formal system specifications.

The overarching goal of the seminar was thus to address the following question: *Is there a universal specification formalism that can be used as a standard language for a variety of modern cyber-physical systems?* To address this question, this seminar was divided into three broad thrusts:

- State-of-the-art in general specification formalisms,
- Domain-specific needs and domain-specific specification formalisms,
- Expressivity, Monitoring Algorithms and Analysis concerns for a specification language.

Outcome of the seminar

The seminar had a total of 37 participants with a mix of research communities including experts (both theoreticians and practitioners) in formal methods, runtime monitoring, machine learning, control theorey, industrial IoT, and biological systems. The seminar focused on the cross-domain challenges in the development of a universal specification formalism that can accommodate for various CPS applications.

The seminar provided an excellent overview of requirements from various application domains that paved the road for identifying common features in a cross-domain specification language. As another outcome of the seminar, we defined as a community the following next steps:

1. Identification of various benchmark problems for monitoring specifications at runtime, and learning specifications from data.
2. Standardizing syntax for expressing time-series data, such as comma separated values (CSV) with a well-defined header file.
3. Creating a public repository containing traces, specifications, models, and pattern libraries.
4. Coordination with RVComp, a runtime verification competition collocated with the Runtime Verification (RV) conference, and possible coordination with SygusComp (Syntax-guided synthesis competition) and SYNTComp (Synthesis competition) to arrange special tracks on learning specifications.
5. Creating a public repository containing standard parsers for variety of specification formalisms such as variants of Signal Temporal Logic.

Sessions

The seminar was organized as a sequence of open discussions on pre-defined topics of interest. Each session had one or two moderators who introduced the topic and one or two scribes who recorded the proceedings of the discussions. The moderators had a short introduction of the topic, identifying the most important sub-topics for open discussion. The discussions were structured in following sessions:

Day 1 State-of-the-art in general specification formalisms

1. Specification languages in digital hardware
2. Tools perspective
3. Overview of declarative specification languages

Days 2 and 3 Domain-specific needs and domain-specific specification formalisms

1. Specifications in automotive systems
2. Specifications in robotics and perception
3. Specifications in Industry 4.0, EDA and mixed signal design
4. Specifications in smart cities
5. Specifications in bioloty
6. Specifications in medical devices
7. Specifications in security

Days 4 Expressivity, monitoring algorithms and analysis concerns

1. Algorithms for specifications: specifications for learning *versus* learning specifications
2. Streaming languages
3. Runtime monitoring
4. Expressivity

Day 5 Next steps and summary of the seminar outcomes

We also organized on Day 1 a session to honor the memory of Oded Maler, one of the co-organizers of this seminar, and who sadly passed away in September 2018.

6.11 The Role of Non-monotonic Reasoning in Future Development of Artificial Intelligence

Organizers: Anthony Hunter, Gabriele Kern-Isberner, Thomas Meyer, and Renata Wassermann

Seminar No. 19072

Date: February 10–15, 2019 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.9.2.73

Participants: Zied Bouraoui, Arina Britz, Giovanni Casini, Célia da Costa Pereira, Luc De Raedt, James P. Delgrande, Thomas Eiter, Eduardo Fermé, Laura Giordano, Andreas Herzig, Anthony Hunter, Gabriele Kern-Isberner, Sebastien Konieczny, Maria Vanina Martinez, Thomas Meyer, Abhaya Nayak, Odile Papini, Marco Ragni, Gavin Rens, Hans Rott, Ken Satoh, Guillermo R. Simari, Ivan José Varzinczak, Serena Villata, Renata Wassermann, Emil Weydert

Nonmonotonic reasoning (NMR) addresses a fundamental problem that classical logic methods in computer science encounter when modelling real-world problems: New information may not only extend previously held knowledge (this would correspond to a monotonic extension) but can drastically change knowledge in that conclusions turn out to be wrong and need to be withdrawn. Nonmonotonic phenomena are present in all areas of our everyday lives mostly due to uncertain and incomplete information, but also due to humans reasoning with restricted ressources; on the other hand, humans do very well in determining relevant contexts of reasoning, so reasoning from incomplete information only may well be on purpose and for sake of efficiency. Nowadays, with computer systems taking on increasingly sophisticated roles in our lives, the need for computational intelligence to be able to also reason in a nonmonotonic way becomes increasingly urgent.

The international Nonmononotonic Reasoning (NMR) workshops have provided a premier specialized forum for researchers in non-monotonic reasoning and related areas since 1984. Over the years, NMR topics and results have been developed in areas such as answer set programming, computational models of argument, and description logics for ontologies. However, research on core topics of NMR has been scattered into different subcommunities that no longer collaborate in depth on a regular basis. As a consequence, much time and effort for solving specific, but in principle similar problems is wasted, general relevance of proposed solutions is overlooked, and general methodological competence is no longer developed to the same degree as ten years ago.

This Perspectives Seminar brought together researchers both from core topics and peripheral areas of NMR, but also attracted researchers from other scientific domains in which recent developments have shown an increased relevance of NMR topics. More precisely, researchers from various subcommunities within computer science and engineering (e.g., artificial intelligence, classical and non-classical logics, machine learning, agent and multiagent systems) met in Dagstuhl, but also researchers from other disciplines like philosophy and psychology contributed to the seminar. The overall goal of this seminar was to reshape NMR as a core methodology for artificial intelligence being able to meet present and future challenges. For AI to progress from pattern recognition and machine learning to broader cognitive reasoning, it needs to have commonsense reasoning, and this in turn calls for a deeper understanding of NMR. So participants of this workshop discussed in what shape NMR would be useful for future AI, and how NMR can be developed for those requirements. We started the seminar with brief survey talks on answer set programming, belief revision, argumentation, argument mining, machine learning, conditional reasoning, description logics, as well as NMR and cognition, and had some technical talks on central topics of NMR afterwards. For the rest of the week, we had working groups on NMR and learning, NMR and cognition, engineering NMR, and commonsense reasoning. We let people freely choose which working groups they wanted to attend each day, which resulted in vivid discussions and a particularly dynamic exchange of ideas. On the last day of the seminar, each working group presented their ideas and future plans, and we closed this seminar with a plenary discussion on the future of NMR. This report shows brief summaries of the presentations and of the results of the working groups.

6.12 Verification and Synthesis of Human-Robot Interaction

Organizers: Rachid Alami, Kerstin I. Eder, Guy Hoffman, and Hadas Kress-Gazit

Seminar No. 19081

Date: February 17–22, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.2.91

Participants: Erika Abraham, Henny Admoni, Rachid Alami, Dejanira Araiza-Illan, Brenna D. Argall, Frank Broz, Maya Cakmak, Alessandro Cimatti, Aurelie Clodic, Kerstin I. Eder, Rüdiger Ehlers, Victor Fernandez Castro, Ivan Gavran, Michael Gienger, Marc Hanheide, Christoffer R, Heckman, Guy Hoffman, Shanee Honig, Felix Ingrand, Nils Jansen, Ross A. Knepper, Hadas Kress-Gazit, Jan Kretinsky, Morteza Lahijanian, Ute Leonards, Shelly Levy-Tzedek, Jamy Jue Li, Daniele Magazzeni, Björn Matthias, Todd Murphey, Laurel Riek, Kristin Yvonne Rozier, Dorsa Sadigh, Maha Salem, Satoru Satake, Jana Tumova

There is a growing trend in robotics moving from industrial robots that work physically separated from people to robots that collaborate and interact with people in the workplace and the home. The field of human-robot interaction (HRI) studies such interactions from the computational, design and social points of view. At the same time, there is growing interest in research regarding the safety, verification and automated synthesis of behaviors for robots and autonomous systems. The fields of formal methods and testing, which focus on verification and synthesis of systems, aim to model systems and define and prove specifications over these systems; in the context of robotics, these techniques take into account the robot dynamics and its interaction with its changing and uncertain environment.

However, a human collaborating with a robot is not just part of the robot's environment, but an autonomous agent with intentions, beliefs, and actions that mesh with those of the robotic agent. This raises new research questions related to verification and synthesis including what appropriate models for human-robot interaction would be; whether and how algorithms for HRI can enable verification; how to take the human into account in automatic synthesis of robotic systems; and what (if any) guarantees can be provided with a human in the loop.

To date, very little work has explored questions of verification, safety guarantees and automated synthesis in the context of Human-Robot Interaction. HRI has modeled humans computationally but not from a verification point of view and without providing guarantees. Furthermore, there are rarely any formal specifications in the computational HRI literature; validated objective metrics for evaluation are also scarce. The verification and synthesis community has mostly focused on the robot's autonomous behavior and its environment, and not paid much attention to the integral presence of the human or the interaction, including the psychological, social, and intentional aspects of human activity.

In this seminar we bring together experts in computational HRI, verification of autonomous systems, formal methods, and cognitive and social psychology to exchange ideas, define research directions, and foster collaborations toward a new theory and practice of verifiable HRI.

6.13 AI for the Social Good

Organizers: Claudia Clopath, Ruben De Winne, Mohammad Emtiyaz Khan, and Tom Schaul

Seminar No. 19082

Date: February 17–22, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.2.111

Participants: Gerald Abila, Hisham Almiraat, Hiromi Arai, Danielle C. M. Belgrave, Fanny Cachat, Claudia Clopath, Bec Connelly, Julien Cornebise, Wilfried de Wever, Ruben De Winne, Daphne Ezer, Tobias Glasmachers, Frank Hutter, Mohammad Emtiyaz Khan, Shakir Mohamed, Frank Mugisha, Mihoko Otake, Mustafa Othman, Angela Picciariello, Julia Proskurnia, Tom Schaul, Kyle Snyder, Yee Whye Teh, Nenad Tomasev

The purpose of Dagstuhl Seminar 19082: AI for the Social Good was to bring together researchers in artificial intelligence (AI) and machine learning (ML) with non-governmental organisations (NGOs) to explore if and how AI and ML could benefit the social good. Indeed, AI and ML have made impressive progress in the last few years. Long-standing challenges like Go have fallen and the technology has entered daily use via the vision, speech or translation capabilities in billions of smartphones. The pace of research progress shows no signs of slowing down, and demand for talent is unprecedented. But as part of a wider AI for Social Good trend, this seminar wanted to contribute to ensuring that the social good does not become an afterthought in the rapid AI and ML evolution, but that society benefits as a whole.

The five-day seminar brought together AI and ML researchers from various universities and industry research labs with representatives from NGOs based in Somalia, Rwanda, Uganda, Belgium, United Kingdom and The Netherlands. These NGOs all pursue various social good goals, such as increasing access to justice for vulnerable people, promoting human rights & protecting human rights defenders, and defeating poverty. On these topics, NGOs have rich domain knowledge, just like they have vast networks with (non-)governmental actors in developing countries. Mostly, NGOs have their finger on the pulse of the challenges that the world & especially its most vulnerable inhabitants are facing today, and will be facing tomorrow. The objective of the seminar was to look at these challenges through an AI and ML lens, to explore if and how these technologies could help NGOs to address these challenges. The motivation was also that collaborations between AI and ML researchers and NGOs could benefit both sides: on the one hand, the new techniques can help with prediction, data analysis, modelling, or decision making. On the other hand, the NGOs' domains contain many non-standard conditions, like missing data, side-effects, or multiple competing objectives, all of which are fascinating research challenges in themselves. And of course, publication impact is substantially enhanced when a method has real-world impact.

The seminar facilitated the exploration of possible collaborations between AI and ML researchers and NGOs through a two-pronged approach. This approach combined high-level talks & discussions on the one hand with a hands-on hackathon on the other hand. High-level talks & discussions focused first on the central concepts and theories in AI and ML and in the NGOs' development work, before diving into specific issues such as privacy & anonymity, data quality, intellectual property, accessibility and ethical issues. These talks and discussions allowed all participants – in a very short timeframe – to reach a sufficient level of understanding of each other's work. This understanding was the basis to then start investigating jointly through a hackathon how AI and ML could help addressing the real-world challenges presented by the NGOs. At the start of the hackathon, an open marketplace-like setting allowed AI and ML researchers and NGOs to find the best match between technological supply and demand. When teams of researchers and NGOs were established, their initial objective was not to start coding, but to define objectives, assess scope and feasibility. Throughout the hackathon, group membership was fluid, as some projects finished early, were deemed out of scope, or needed to wait for data. Some groups managed to build a viable initial prototype, others established the seeds for future collaborations, and a few were proposed as full summer projects within the "Data Science of Social Good summer school". The projects' aims were diverse. They included better seeds for farmers, modelling cognitive age and decline, scalable legal assistance and scalable citizen feedback. As a result of the hackathon, all NGOs could take concrete results home – some to build on further, some as mature solutions.

Finally, a result of the seminar that is relevant for the entire AI for Social Good community are the ten key challenges for AI for Social Good initiatives that participants identified:

1. the importance of deep, long-term partnerships,
2. clear and well-defined goals and use cases,
3. bias towards simpler solutions,
4. data readiness,
5. setting expectations with regards to both impact and the pace at which technology can be applied,
6. ensuring privacy and security of data,
7. inclusivity and ethics of the applications,
8. factoring in the limitations of both communities,
9. challenges in overcoming the barriers to NGOs utilising the potential of AI/ML, and
10. the relative cost of AI/ML for social good.

6.14 Beyond-Planar Graphs: Combinatorics, Models and Algorithms

Organizers: Seok-Hee Hong, Michael Kaufmann, János Pach, and Csaba D. Tóth

Seminar No. 19092

Date: February 24– March 1, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.2.123

Participants: Eyal Ackerman, Carlos Alegría Galicia, Patrizio Angelini, Michael Bekos, María del Pilar Cano Vila, Giordano Da Lozzo, Vida Dujmović, Peter Eades, Stefan Felsner, Henry Förster, Fabrizio Frati, Radoslav Fulek, Martin Gronemann, Michael Hoffmann, Seok-Hee Hong, Michael Kaufmann, Balázs Keszegh, Linda Kleist, Stephen G. Kobourov, Giuseppe Liotta, Anna Lubiw, Tamara Mchedlidze, Fabrizio Montecchiani, Martin Nöllenburg, Yoshio Okamoto, János Pach, Maurizio Patrignani, Sergey Pupyrev, Chrysanthi Raftopoulou, Günter Rote, Ignaz Rutter, Leonie Ryvkin, Csaba D. Tóth, Géza Tóth, Torsten Ueckerdt, Alexander Wolff

Most big data sets are relational, containing a set of objects and relations between the objects. This is commonly modeled by graphs, with the objects as the vertices and the relations as the edges. A great deal is known about the structure and properties of special types of graphs, in particular *planar graphs* which are fundamental for both Graph Theory, Graph Algorithms and Automatic Layout. Structural properties of planar graphs can often be expressed, for example, in terms of excluded minors, low density, and small separators. These properties lead to efficient algorithms; consequently a number of fundamental algorithms for planar graphs have been discovered. As many of the characteristic properties of planar graphs have been generalized (e.g., graph minor theory, topological obstructions, χ-boundedness), these algorithms also extend in various directions to broad families of graphs.

Typical real world graphs, such as social networks and biological networks, are *nonplanar*. In particular, the class of scale-free networks, which can be used to model web-graphs, social networks and many kinds of biological networks, are sparse nonplanar graphs, with globally sparse and locally dense structure. To analyze and visualize such real world networks, we need to formulate and solve fundamental mathematical and algorithmic research questions on *sparse nonplanar* graphs. Sparsity, in most cases, is explained by properties that generalize those of planar graphs: in terms of topological obstructions or forbidden intersection patterns among the edges. These are called *beyond-planar graphs*. Important beyond-planar graph classes include the following:

- *k-planar graphs*: graphs that can be drawn with at most k crossings per edge;
- *k-quasi-planar graphs*: graphs which can be drawn without k mutually crossing edges;
- *k-gap-planar graphs*: graphs that admit a drawing in which each crossing is assigned to one of the two involved edges and each edge is assigned at most k of its crossings;
- *RAC (Right Angle Crossing) graphs*: graphs that have straight-line drawings in which any two crossing edges meet in a right angle;
- *bar k-visibility graphs*: graphs whose vertices are represented as horizontal segments (bars) and edges are represented as vertical lines connecting bars, intersecting at most k bars;
- *fan-crossing-free graphs*: graphs which can be drawn without fan-crossings; and
- *fan-planar graphs*: graphs which can be drawn such that every edge is crossed only by pairwise adjacent edges (fans).

Compared to the first edition of the seminar, we planned to focus more on aspects of computational geometry. Therefore, we included one new organizer as well as some more participants from this field.

Thirty-six participants met on Sunday afternoon for a first informal get-together and reunion since the last workshop which took place more than two years ago. From that event, the four working groups nearly all have completed and published subsequent work. We decided to build on the achievements of the previous meeting and scheduled short talks recalling the previous seminar's results. On Monday afternoon, we held an engaging open problems session and formed new working groups. Notably, this time, more problems related to computational geometry as well as questions from combinatorics have been proposed. Open problems included questions about the combinatorial structures (e.g, book thickness, queue number), the topology (e.g., simultaneous embeddability, gap planarity, quasi-quasiplanarity), the geometric representations (e.g., representations on few segments or arcs), and applications (e.g., manipulation of graph drawings by untangling operations) of beyond-planar graphs.

In the opening session of every morning, we have drawn inspiration from additional talks, fresh conference contributions on related topics (see abstracts). An impressive session on the last day was devoted to progress reports that included plans for publications and follow-up projects among researchers that would have been highly unlikely without this seminar. From our personal impression and the feedback of the participants, the seminar has initiated collaboration and lead to new ideas and directions.

We thank all the people from Schloss Dagstuhl for providing a positive environment and hope to repeat this seminar, possibly with some new focus, for a third time.

6.15 Analysis, Design, and Control of Predictable Interconnected Systems

Organizers: Kunal Agrawal, Enrico Bini, and Giovanni Stea

Seminar No. 19101

Date: March 3–8, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.1

Participants: Kunal Agrawal, Luis Almeida, Sanjoy K. Baruah, Enrico Bini, Steffen Bondorf, Anne Bouillard, Marc Boyer, Peter Buchholz, Giacomo Como, Arvind Easwaran, Pontus Ekberg, Johan Eker, Markus Fidler, Laura Galli, Fabien Geyer, Sathish Gopalakrishnan, James Gross, Chadlia Jerad, Yuming Jiang, Li Jing, Pratyush Kumar, Kai Lampka, Jean-Yves Le Boudec, Edward A. Lee, Jörg Liebeherr, Martina Maggio, Ahlem Mifdaoui, Victor Millnert, Geoffrey Nelissen, Paul Nikolaus, Amr Rizk, Ketan Savla, Giovanni Stea, Niklas Ueter, Tongtong Wang, Kui Wu, Kecheng Yang, Jiayi Zhang, Marco Zimmerling

Computing capacity surronds our environment more and more. The exploitation of this diffused infrastructure poses a number of challenges. A notable one is the guarantee of predictable performance. Many modern computing platforms require high degrees of predictability in their timing characteristics – for instance, in avionics and automotive applications, it is essential that the tasks complete in a timely manner in order to take appropriate actions in response to a developing situation. With increasing complexity and heterogeneity in functionality required by time-sensitive applications, there is an increasing need for designing distributed interconnected platforms that respond to time-varying requests in a highly predictable way.

The world of networking is undergoing an analogous transformation. There is a major shift towards smarter and more autonomous networks, the so-called self-driving networks. The goal is to mimic the success that cloud-computing techniques and concepts had on the transformation of the IT infrastructure. The latter made it possible to allow physical resources, such as computing and storage nodes, to be shared among users through the use virtual resources (Network Function Virtualization). Virtualization of network functions offers an efficient way to meet dynamic user demands in a cost-effective manner. However, the guarantee of predictability in NFV is, to date, an open problem.

With increasing complexity and heterogeneity in functionality required by timing-sensitive applications, there is need for designing distributed and interconnected platforms that provide high predictability. On the one hand, with the advent of Cyber-physical systems and Industry 4.0, real-time systems are becoming more and more networked systems. On the other hand, networking scenarios where a-priori quantifiable guarantees on the worst-case traversal time are required as a prerequisite for a correct application-level computation are becoming commonplace. Therefore, it is necessary to develop methods to ensure compliance with end-to-end deadlines in distributed systems where *both* computation and communication are involved.

Dagstuhl seminar "Analysis, design and control of predictable interconnected systems" gathered more than 40 researchers from four continents, with a good balance of seniority and gender. These researchers came from different communities, including Network Calculus, Real-time Systems, Control Theory, Performance Evaluation and Data-flow, which have been discussing similar problems in the recent past, sometimes also using similar methodologies, but different notations and hypotheses, and focusing on different characteristics. The aim of the seminar was to foster cross-fertilization and inter-community networking, using new, contemporary problems which are interesting for and discussed within more than one community. The seminar hosted 26 talks. Some of these were "background-levelling" talks given by well-recognized senior experts in the resepective communities, with the aim of crossing the lexicon and notation gap between communities. Other talks, building on the shared background provided by the former, discussed interesting novel problems and promising application areas. The open, schedule-as-you-go format for the talks and the time left open for free networking activities fostered an open environment. The general conclusion, which can be gathered by the (mostly enthusiastic) comments in the inal survey, is that the involved communities have more in common with each other than the attendees thought, and cross-fertilization is necessary to tackle new problems in a holistic approach.

6.16 3D Morphable Models

Organizers: Bernhard Egger, William Smith, Christian Theobalt, and Thomas Vetter

Seminar No. 19102

Date: March 3–8, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.16

Participants: Thabo Beeler, Florian Bernard, Michael J. Black, Volker Blanz, Timo Bolkart, Bernhard Egger, Victoria Fernandez Abrevaya, Patrik Huber, Ron Kimmel, Tatsuro Koizumi, Adam Edward Kortylewski, Yeara Kozlov, Andreas Morel-Forster, Nick Pears, Gerard Pons-Moll, Emanuele Rodolà, Sami Romdhani, Javier Romero, Shunsuke Saito, William Smith, Ayush Tewari, Christian Theobalt, Thomas Vetter, Stefanie Wuhrer, Michael Zollhöfer, Silvia Zuffi

A total of 45 people was invited to this seminar in the first round of invitations. The seminar was fully booked after the first round and 26 researchers from academia and industry participated in the seminar. 21 researchers presented their work in around 15-30 minutes presentations, an abstract of each presentation is included in this report. Besides those presentations participants where presenting their shared data and software in a specific slot. We collected this information in a list of shared resources which we made publicly available[48]. This overview and exchange was one of the aims we had initially in mind when organizing the workshop. In the beginning of the workshop we collected ideas for discussions in our flexible sessions, those ideas are also contained in this report. We then structured the seminar fixing the topics of discussion for the flexible sessions. The summaries of those discussions are also contained in this report. One slot was reserved for a joint group discussion on upcoming ethical concerns on the methods we are developing. This interesting and well organized discussion was an initiative from the participants and not foreseen by the organizers. Another bigger discussion was around the topic of how to compare different approaches and how to establish a benchmark. We did not completely converge on a final solution but we identified currently available benchmarks and we discussed how a gold-standard benchmark would look like. Another aim of the workshop was to initiate an edited book or a survey paper with broad support. Arising from the workshop a group of 13 junior and senior researchers started to work on a joint survey and perspective paper on 20 years of Morphable Face Models. Discussions and presentations were followed by vivid discussions on current challenges and future research directions. To future nurture the ideas of the seminar we started a google group for discussions, sharing news and exchanging students [49].

The group would like to meet again at Dagstuhl in 2022. The program was more dense than expected and we would like to have more time for discussions in groups after a set of talks. We would like to highlight 5 main discussion points:

- To what degree of detail we need to model in 3D and physically adequate, what can we learn from semi-supervised or unsupervised 2D data?
- Is the model depending on the application or is there a golden standard model that is able to fit all applications?
- The current revolution of deep learning in computer vision enables a lot of novel strategies and speeds up the models, however, other challenges in modeling, synthesis and inverse rendering remain and new deep learning specific challenges are introduced.
- What are the ethical implications of the models and systems we are building?
- How will the field develop in the next 20 years? Which challenges should we focus on?

We started the seminar with a short introduction of everybody. The homework was to introduce themself with at most one slide and prepare one important question, challenge or goal you would like to discuss during the seminar.

- **Thabo Beeler**: Non-Linear Morphable Models. How to get off Model in a meaningful way?
- **Florian Bernard**: Deeper integration of models of human knowledge and algorithms into learning systems. What are potential perspectives? How to best approach this?
- **Michael J. Black**: What's next? Increasing realism? Deep representations? Something else?
- **Volker Blanz:** Expressive model also reproduces non-face

[48] https://github.com/3d-morphable-models/curated-list-of-awesome-3D-Morphable-Model-software-and-data
[49] https://groups.google.com/forum/#!forum/3d-morphable-models

structures! How to discriminate between face and non-face? Future: better regularization, rely on trained regressors, recognize glasses ?
- **Bernhard Egger:** What to model? What to learn?
- **Victoria Fernandez Abrevaya:** How far are we from closing the gap between high-quality and low-quality capture devices, and can we use 3DMM for this?
- **Patrik Huber:** What is missing to reliably reconstruct realistic 3D faces from mostly uncontrolled 2D footage?
- **Ron Kimmel:** Geometry is the art of finding the "right" parametrization. Deep Learning is a technology that exploits convenient parametric spaces (CNN) for classification. Any hope for unification? Is translating geometry into algebra the answer?
- **Tatsuro Koizumi:** How to evaluate and assure the robustness of neural network-based reconstruction? How to improve the stability of self-supervised training?
- **Adam Edward Kortylewski:** Can we resolve the limitations of Deep Learning with Generative Object Models?
- **Yeara Kozlov:** Can physically based face modeling be replaced by machine learning?
- **Andreas Morel-Forster:** Fast posterior estimation – A contradiction?
- **Nick Pears:** How to build deeper, wider models?
- **Gerard Pons-Moll:** Is the Euclidean 3D space the right space to model humans, clothing and hair?
- **Emanuele Rodolà:** Can we make inverse spectral geometry useful in practice?
- **Sami Romdhani:** How to combine Deep Learning and 3D Equations to generate images?
- **Javier Romero:** How can Deep Nets learn from unstructured, uncalibrated views?
- **Shunsuke Saito:** Is there an unified representation to represent digital human without explicitly having prior for each component?
- **William Smith:** Self-supervision: holy grail or just re-discovering gradient descent-based analysis-by-synthesis? How do we make sure the gradients of our losses are really useful (Appearance loss: meaningless when far from good solution, Landmark loss: ambiguous (and not self-supervised), Rasterization: not differentiable)?
- **Ayush Tewari:** How can we build high quality 3D morphable models from 2D data?
- **Christian Theobalt:** Can we build a 4D Real World Reconstruction Loop? Ethical, Privacy, Security Questions of Parametric/Morphable Model Building and Reconstruction Algorithms
- **Thomas Vetter:** Did we learn much about this optimization problem (inverse rendering)?
- **Stefanie Wuhrer:** How to effectively learn parametric human models from captured data using minimal supervision?
- **Michael Zollhöfer:** What is the best representation for deep learning-based 3D reconstruction and image synthesis?
- **Silvia Zuffi:** How to model skin dynamics from video?

After the individual introductions, we discussed those ideas in discussion groups to identify points to discuss during the seminar. The following list is the unfiltered result of our brainstorming on open questions and challenges.

- Where to spend the next 20 years? Perfection: finer detail? Move it: Movement, new representation, new goals, new data? Break it: hair, clothing, new representation, new goals, new data?
- Why aren't we focusing on fixing the obvious errors?
- Optimization: Why aren't we doing more to understand our objective function and adopt the algorithms?
- How to predict distributions instead of point estimates?
- How much detail to model vs. overfitting?
- How to evaluate Photorealism?
- Should vision people be more aware of graphics standard for photorealism?
- Is it important to understand?
- Do we need correspondences to build 3D models and predictions?
- How to learn 3D from 2D?
- How to adapt models over time (without calibration)?
- How to deal with multi-view and video in CNNs?
- Which courses/skills are required?
- Use for society?
- What to leave for industry?
- What is the role of academia within industry (collaboration vs. isolation)?
- Representations (beyond triangle meshes) to deal with category discontinuities, e.g. smooth surface vs. hair
- Evaluation of shape and appearance reconstruction
- Connections between deep learning and parametric models
- Role of axiomic models in learning
- Comparability: Benchmark and metrics
- Future prediction of motion
- Self-supervision
- Differentiable inverse rendering

6.17 Theoretical Foundations of Storage Systems

Organizers: Martin Farach-Colton, Inge Li Gørtz, Rob Johnson, and Donald E. Porter

Seminar No. 19111

Date: March 10–15, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.39

Participants: Michael A. Bender, Ioana Bercea, Jonathan Berry, Philip Bille, Timo Bingmann, Alexander Conway, Guy Even, Martin Farach-Colton, Jeremy Fineman, Johannes Fischer, Pawel Gawrychowski, Seth Gilbert, Inge Li Gørtz, Magnús M. Halldórsson, William Jannen, Rob Johnson, Tomasz Kociumaka, Geoff Kuenning, Bradley C. Kuszmaul, William Kuszmaul, Simon Mauras, Samuel McCauley, Ulrich Carsten Meyer, Miguel A. Mosteiro, Ian Munro, Sam H. Noh, Prashant Pandey, Nikos Parotsidis, Cynthia A. Phillips, Solon Pissis, Donald E. Porter, Simon J. Puglisi, Tom Ridge, Eva Rotenberg, Siddhartha Sen, Francesco Silvestri, Shikha Singh, Meng-Tsung Tsai, Przemyslaw Uznanski, Janet Vorobyeva, Gala Yadgar

Storage systems, including databases and file systems, are at the heart of all large data applications. Recently, some storage systems have incorporated theoretical advances in data organization techniques, with substantial improvements in performance. However, there is little contact between the systems designers who build storage systems and theoreticians who design new ways to organize data. This Seminar workd to bridge this gap, to the benefit of both communities and to improve the design of all storage systems.

External-memory algorithms are those where the data is too large to fit in memory, and hence needs to be stored on disk and accessed using I/O. Algorithmic analysis of such algorithms therefore focuses on the number of I/Os needed to complete a computation, rather than the number of instructions. This is because an I/O can be many orders of magnitude slower than a machine instruction and therefore I/Os can be the bottleneck in such computations.

The theoretical analysis of such external-memory algorithms has produced many exciting results in the last two decades. Many of these are directly relevant to practical applications, but only a few have made the leap to deployment. This Seminar brought together theoreticians, who have an extensive understanding of the state of the art in external-memory data structures, and storage systems researchers and practitioners, who understand the details of the problems that need to be solved.

Specific questions that were addressed in the workshop include the following:

- How can we use the huge improvements in string data structures to improve storage systems that manipulate strings? Many data structures, such as LSMs and B^ϵ-trees, rely heavily on the assumption that keys are indivisible and small.
- How can we use new multi-dimensional data indexes in working systems?
- Many indexes suffer from fragmentation. Are there data structure improvements that would allow efficient storage on disks that are nearly full? Currently, disks are kept only a fraction full because the performance of existing data structures decays dramatically as the disk fills. This suggests another problem:
- How can theoretical models be improved to capture such issues as:
  - full disk: The external memory model, the disk is of unbounded size.
  - sequential access: Both hard disks and SSDs require very large blocks of sequential I/O to capture a large fraction of bandwidth. The external-memory model assumes that disks are random access.
- Concurrency: Can data structures be designed to exploit memory locality on disk while maintaining concurrency in RAM?

The storage system world is in a ferment as new hardware becomes available. Now is the time to establish deep partnerships across disciplines in computer science to solve some of the most pressing big data infrastructure problems.

6.18 Engineering Reliable Multiagent Systems

Organizers: Jürgen Dix, Brian Logan, and Michael Winikoff

Seminar No. 19112

Date: March 10–15, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.52

Participants: Tobias Ahlbrecht, Stefano Albrecht, Natasha Alechina, Rem Collier, Mehdi Dastani, Louise A. Dennis, Frank Dignum, Virginia Dignum, Jürgen Dix, Niklas Fiekas, Michael Fisher, Koen V. Hindriks, Alexander Birch Jensen, Malte S. Kließ, Yves Lesperance, Brian Logan, Viviana Mascardi, Ann Nowé, Alessandro Ricci, Kristin Yvonne Rozier, Holger Schlingloff, Marija Slavkovik, Kagan Tumer, Michael Winikoff, Neil Yorke-Smith

The multi-disciplinary workshop on Reliable Multiagent Systems attracted 26 leading international scholars from different research fields, incuding theoretical computer science, engineering multiagent systems, machine learning and ethics in artificial intelligence.

This seminar can be seen as a first step to establish a new research agenda for engineering reliable autonomous systems: clarifying the problem, its properties, and their implications for solutions.

In order to move towards a true cross-community research agenda for addressing the overarching challenge of engineering reliable autonomous systems we have chosen a slightly different organization than usual: the seminar was comprised of (short) talks (days 1 and 2), and extensive discussions and dedicated group work (days 3-5).

The first two days were opened by two longer (45 minutes each) tutorials, followed by short "teaser talks" ($10 + 5$ minutes) related to the main topic of *reliable multiagent systems*. Almost all participants gave their view of the topic and highlighted possible contributions. The talks were meant to be less "conference-style", and more inspiring and thought-provoking.

At the end of the second day, we established four working groups to dive deeper into the following questions:

1. What **(detailed) process** can be used to identify properties that a particular reliable autonomous system or MAS needs to satisfy?
2. How can we engineer reliable autonomous systems that include **learning**?
3. How can we engineer reliable autonomous systems that include **human-machine interaction** (including human-software teamwork)?
4. How can we engineer reliable autonomous systems comprising **multiple agents** (considering teamwork, collaboration, competitiveness, swarms, ...)?

The groups met on Wednesday and Thursday for extensive discussions and reported back intermediate results in plenary sessions. Participants were encouraged to move between groups to enrich them with their expertise. The seminar concluded on Friday morning with a general discussion where all groups presented their results.

We summarise below the key results from the four discussion groups.

Identifying properties: This group discussed the challenge of identifying requirement properties to be verified. It focused in particular on autonomous systems that replace humans in domains that are subject to regulation, since these are most likely to require and benefit from formal verification.

The group articulated the following arguments:

- That autonomous systems be viewed in terms of three layers: a continuous control layer at the bottom, a "regulatory" layer in the middle, and an "ethical" layer at the top. The distinction between the regulatory and ethical layers are that the former deals with the expected normal behaviour (e.g. following the standard rules of the domain), whereas the latter deals with reasoning in situations where the rules need to be broken. For example, breaking a road rule given appropriate justification.
- That for these sorts of systems we can derive verification properties by considering the licencing that is used for humans and how human skills and capabilities are assessed, as well as relevant human capabilities that are assumed, and relevant laws and regulations. The group sketched a high-level process for identifying requirement properties by considering these factors.

The group considered a range of domains, for each one showing how these points would apply.

These ideas were developed into a draft paper during the workshop, and work on this paper has continued subsequently.

Learning in reliable autonomous system: The second group was concerned with methods for engineering reliable autonomous systems that involve learning.

The notion of sufficient reliability varies from domain to domain. For example, in planning of telecommunication networks there are simulators that are trusted to be a good model of reality. Hence the simulation rules could be used for formal verification. In other domains, such as autonomous driving, there is no established trusted model of reality.

Assuming a formal model exists and safety properties can be formulated with temporal logic, there are still remaining challenges: complex models with a large state space and hybrid continuous and discrete behavior can make formal verification intractable, especially when the learned policies are equally complex. On the other hand learning methods (e.g. reinforcement learning) often "discover" key strategies that do not depend on all details of the system. The group discussed ideas for abstracting/discretizing transition systems based on learned policies.

Human-machine interaction in reliable autonomous systems: The third group focused on how to engineer reliable human-agent interaction.

For that, the first step was to carve out what it means for human-machine communication to be reliable. Values and norms are definitely involved. Drawing from human communication, being truthful, up-to-date with relevant knowledge and honouring commitments are major parts. Another important building block is transparency: is it always clear, which values are in play, what the agent's purpose is, or what happens with the collected data? The desired result would be reliable outcomes, e.g. through reliably following a protocol, effective communication and getting to a shared understanding. A number of tools and methods to achieve this were identified: stakeholder analysis, plan patterns/activity diagrams, interaction design patterns, appropriate human training, and explainability (i.e. explainable AI) were among the most prominent engineering solutions. This group also concluded their discussions early and distributed themselves among the other groups after that.

Multiple agents in reliable autonomous systems: The fourth group focussed on the challenges of ensuring reliability in systems comprising multiple, possibly heterogeneous, agents interacting in complex ways.

A number of issues emerged from the discussion, including what does it mean for a multiagent system to be "collectively reliable", and what is the relationship between the reliability or otherwise of individual agents and the coordination mechanisms through which they interact, and the collective reliability of the system of a whole. These issues were broken down into more specific engineering challenges, including which languages should be used to express collective reliability properties (which is closely related to the discussion of the first group) and how such properties should be verified, how to engineer reliable coordination mechanisms when we have only partial access to the states of agents, how to decompose and/or distribute the monitoring and control of individual agents (and associated failure recovery) necessary for reliable coordination, how to do all of the above in systems where agents learn (closely related to the discussion of the second group), and, finally, how to allocate responsibility to individual agents when behaviour is not reliable.

A more detailed research agenda for engineering reliable multiagent systems is in preparation, which we plan to publish as a "position paper" in a journal special issue arising from the work at the workshop.

6.19 Computational Complexity of Discrete Problems

Organizers: Anna Gál, Rahul Santhanam, and Till Tantau

Seminar No. 19121

Date: March 17–22, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.64

Participants: Eric Allender, Paul Beame, Harry Buhrman, Igor Carboni Oliveira, Katrin Casel, Amit Chakrabarti, Sourav Chakraborty, Arkadev Chattopadhyay, Gil Cohen, Anindya De, Lukáš Folwarczný, Lance Fortnow, Anna Gál, Alexander Golovnev, Kristoffer Arnsfelt Hansen, Prahladh Harsha, Johan Hastad, Valentine Kabanets, Michael Kapralov, Mathew Katzman, Antonina Kolokolova, Swastik Kopparty, Michal Koucký, Matthias Krause, Meena Mahajan, Or Meir, Jakob Nordström, Ramamohan Paturi, Pavel Pudlák, Rüdiger Reischuk, Michael E. Saks, Rahul Santhanam, Ronen Shaltiel, Amnon Ta-Shma, Avishay Tal, Till Tantau, Thomas Thierauf, Jacobo Torán, Ben Lee Volk, Omri Weinstein

Computational complexity theory is the study of computation under bounded resources, and the tradeoffs thereof offered by specific problems and classes of problems in various computational models. Such resources include time and space for classical computation, randomnesss, non-determinism, and oracles for more advanced uniform machines, size/advice for circuits/non-uniform computation, interaction for communication protocols, length and depth for proof complexity, and much more. The goals of work in this field are not only to discover and improve these tradeoffs, but ideally to find tight lower bounds to match the solutions that have been found, and use such results in one of the models to inform results in the others. Despite decades of work on these problems, the answers to many foundational questions (such as **P** vs **NP**, **P** vs **BPP**, **NP** vs **co-NP**) still remain out of reach.

For the 2019 instalment of the seminar series *Computational Complexity of Discrete Problems* – which evolved out of the seminar series *Complexity of Boolean Functions* that dates back to the founding of Dagstuhl – Anna Gál, Oded Regev, Rahul Santhanam, and Till Tantau invited 40 participants to Dagstuhl to work towards discovering new results in the field. The seminar started with the assembly of a large "graph of interests" that allowed the participants both to present their own research interests and to see how these align with the other present researchers. The bulk of the research work was then done in the form of, on the one hand, talks in the morning and late afternoon and, on the other hand, break-out sessions and small discussions in the afternoon by smaller groups.

A distinguishing feature of the seminar talks were the lively discussions during and after the talk: given the often highly abstract and specialized topics presented by the experts in the field, lively discussions are by no means a given and they proved to be both rewarding and helpful for all participants. In the informal afternoon sessions, smaller groups of researchers had ample time to tackle the open problems of the field; with some discussions still going on after midnight. Two events – the traditional Wednesday hike and the traditional wine-and-cheese party on Thursday – allowed everyone well-earned breaks from doing research on computational complexity.

The range of topics covered by the participants during the seminar was broad and included derandomization, lower bounds for specific problems, communication complexity, complexity classes, graph algorithms, learning theory, coding theory, and proof complexity. Specific selected results presented throughout include:

- A proof that the Log-Approximate-Rank Conjecture is false, yielding the first exponential gap between the logarithm of the approximate rank and randomized communication complexity for total functions.
- An oracle separation of **BQP** and the polynomial hierarchy, showing a strong converse to the Bennett et al. oracle relative to which **BQP** cannot solve **NP**-complete problems in sub-exponential time.
- Improved lower bounds for the Minimum Circuit Size Problem, including
  - $\texttt{MCSP} \notin \mathbf{AC^0[p]}$,
  - $\texttt{MCSP}$ requires $N^{3-o(1)}$-size de Morgan formulas,
  - $\texttt{MCSP}$ requires $N^{2-o(1)}$-size general formulas,
  - $\texttt{MCSP}$ requires $2^{\Omega(N^{1/d+2.01})}$-size depth-d $\mathbf{AC^0}$ circuits,

  where the first result is achieved by showing $\texttt{MCSP}$ can solve the coin problem and the others using properties of local pseudorandom generators.

Open problems were posed by Amit Chakrabarty, Alexander Golovnev, Or Meir, and Omri Weinstein.

The organizers, Anna Gál, Oded Regev, Rahul Santhanam, and Till Tantau, would like to thank all participants at this point for the many contributions they made, but we would also like to especially thank the Dagstuhl staff for doing – as always – an excellent job and helping with organizational matters and with making everyone feel welcome.

6.20 Algorithmic Problems in Group Theory

Organizers: Volker Diekert, Olga Kharlampovich, Markus Lohrey, and Alexei Myasnikov

Seminar No. 19131

Date: March 24–29, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.83

Participants: Yago Antolin, Laurent Bartholdi, Montserrat Casals-Ruiz, Laura Ciobanu, Jordi Delgado Rodriguez, Volker Diekert, Andrew Duncan, Bettina Eick, Murray Elder, Michal Ferov, Michael Figelius, Moses Ganardi, Albert Garreta Fontelles, Susan Hermiller, Artur Jez, Ilya Kapovich, Ilya Kazachkov, Olga Kharlampovic, Manfred Kufleitner, Markus Lohrey, Alexei Myasnikov, Volodymyr Nekrashevych, Gretchen Ostheimer, Igor Potapov, Timothy Riley, Paul E. Schupp, Géraud Sénizergues, Vladimir Shpilrain, Rachel Skipper, Tatiana Smirnova-Nagnibeda, Benjamin Steinberg, Howard Straubing, Svetla Vassileva, Alina Vdovina, Enric Ventura Capell, Pascal Weil, Armin Weiß, Georg Zetzsche

The field of combinatorial group theory, a part of abstract algebra, is tightly linked to computational problems from its early days. Already in 1911, i.e., 25 years before Turing's work on the halting problem appeared, Max Dehn introduced and investigated three fundamental group theoretical decision problems: the *word problem*, the *conjugacy problem*, and the *isomorphism problem*. Dehn's problems had a strong influence on the development of modern theoretical computer science. It took more that 40 years before the work of Novikov, Boone, Adjan, and Rabin showed the undecidability of Dehn's decision problems in the class of finitely presented groups. Despite these negative results, for many groups the word problem turned out to be decidable in many important classes of groups. With the rise of complexity theory in the 1960's, also the computational complexity of group theoretic problems moved into the focus of research. From the very beginning, this field attracted researchers from mathematics as well as computer science. Using algorithmic techniques from complexity theory, researchers were able to exhibit highly efficient algorithms for groups, where initially only pure decidability results have been known. A milestone in this context is Lipton and Zalcstein's logspace algorithm for the word problem of finitely generated linear groups. This was the first result putting the word problem for an important class of groups into a complexity class below polynomial time. In the last 10 years, researchers pushed the limits further towards small circuit complexity classes. In particular the class $\mathbf{TC}^0$ turned out be very important in this context. Despite its limited computational power many important group theoretical problems problems were shown to be in $\mathbf{TC}^0$.

Complexity theoretical questions are not the only area where we have seen fruitful interactions between group theory and theoretical computer science in recent years. Other examples can be found in automata theory, data compression, model theory, and reachability problems for infinite state systems. The following paragraphs put some of the seminar talks into the context of these topics and mentions some of the open problems that were discussed during the seminar.

Groups and circuit complexity. Howard Straubing gave an excellent survey on circuit complexity that was particularly addressed to non-experts in complexity theory. Barrington's famous result according to which the word problem for every finite non-solvable groups is hard for $\mathbf{NC}^1$ was explained and several important results centered around the circuit complexity class $\mathbf{TC}^0$ were surveyed. In recent years, $\mathbf{TC}^0$ turned out to be the right class for characterizing the complexity of several group theoretical problems. Two seminar talks presented further examples of group theory problems in $\mathbf{TC}^0$: Armin Weiß gave a talk about the power word problem which is a succinct version of the classical word problem, where powers g^n of group elements with binary encoded integer exponents n are allowed in the input. Despite this succinctness, several power word problems (e.g. for nilpotent groups and certain wreath products of finitely generated abelian groups) can be still solved in $\mathbf{TC}^0$. Moses Ganardi talked on the knapsack problem for finitely generated groups which asks for the solvability of certain exponent equations over a group. Among other results he gave a simple proof showing that the knapsack problem for unary encoded integers is in $\mathbf{TC}^0$. This result generalizes to all finitely generated abelian groups.

Several promising open problems related to the circuit complexity of group theoretical problems were discussed in the open problem sessions: The above mentioned result of Barrington on finite non-solvable groups motivates the question whether the word problem for every finitely generated solvable group is $\mathbf{NC}^1$-hard. Also finding new classes of infinite groups with a word problem in $\mathbf{TC}^0$ is an open research problem that was intensively discussed during the seminar. So far, it is known that solvable linear groups have a word problem in $\mathbf{TC}^0$ and that the class of groups with a word problem in $\mathbf{TC}^0$ is closed under wreath products.

Compression techniques in group theory. Compression techniques turned out to be an important tool for obtaining efficient algorithms in group theory. The general philosophy is trying to avoid storing extremely long words by computing on a compressed representation of these words. This led to the formulation of several succinct versions of classical group theoretical problems, where the group elements in the input are given a succinct version. The power word problem that was introduced by Armin Weiß (see the previous paragraph) is such a succinct problem. The main result of Armin Weiß' talk was an efficient reduction of the power word problem for a free group to the (ordinary) word problem of a free group. It is open whether similar reductions also exist for right-angeled Artin groups and hyperbolic groups.

In the context of solving equations over groups and monoids, the so-called recompression technique led to several important results in recent years. Arthur Jeż (the inventor of this technique) gave a talk on recompression and outlined his non-deterministic linear time algorithm for solvability of word equations. Ciobanu and Elder presented their recent work on equations in hyperbolic groups where they use recompression in order to show that the set of all solutions for a system of equations over a hyperbolic group is an EDT0L language.

Groups and model theory. This research area directly relates to the previous paragraph. The goal is to understand the first-order theory of groups. Of particular interest is the Diophantine theory. Decidability of the Diophantine theory means that one can decide whether a boolean combination of word equations has a solution. Olga Kharlampovich gave a talk about Diophantine theories of metabelian groups. She proved decidability for several important metabelian groups: Baumslag-Solitar groups $\mathrm{BS}(1,n)$ and wreath products $\mathbb{Z} \wr \mathbb{Z}$ and $\mathbb{Z}_n \wr \mathbb{Z}$. Albert Garreta continued this topic and talked about Diophantine theories of solvable groups. He presented a large class of solvable groups (containing for instance all finitely generated non-virtually abelian nilpotent groups and all polycyclic groups that are not virtually metabelian) for which the Diophantine theory is at least as hard as the Diophantine theory of a suitable ring of algebraic integers. This leads to the conjecture that for each member of his family the Diophantine theory is undecidable.

Montserrat Casals-Ruiz talked on the positive theory of groups acting on trees. The positive theory of a group contains all negation-free statements from the full first-order theory. Montserrat Casals-Ruiz proved that many natural examples of groups acting on trees have the same positive theory as a free group of rank two. Ilya Kazachkov presented new results on the full first-order theory of free products and, more generally, graph products of groups. He showed that under certain conditions, elementary equivalent free products (meaning that their first-order theories coincide) must have elementary equivalent factors.

Groups and automata. Besides complexity of algorithmic problems, a very interesting connection between group theory and theoretical computer science is provided by automata theory, using the very flexible and algorithmically efficient finite state automata to somehow describe an infinite group. This led to the development of automatic groups and automaton groups. Automaton groups are a subclass of so-called self-similar groups. Laurent Bartholdi gave a talk on algorithmic results on self-similar groups and outlined the proof of a recent breakthrough result of Bartholdi and Mitrofanov stating that there exist self-similar groups with an undecidable word problem. For the particular case of automaton groups the word problem belongs to $\mathbf{PSPACE}$. The question whether there exist automaton groups with a $\mathbf{PSPACE}$-complete word problem was intensively discussed during the seminar. Recently, as a direct outcome of the seminar, an automaton group with this property was constructed by Jan Philipp Wächter and Armin Weiß [An automaton group with PSPACE-complete word problem, arXiv, 2019. https://arxiv.org/abs/1906.03424]. Volodia Nekrashevych presented in his talk a generalization of automaton groups based on non-deterministic synchronous automata-transducers and discussed their algorithmic properties and relationship to dynamical systems.

Reachability problems. The study of reachability problems for matrix semigroups has a long tradition in theoretical computer science. Formulated in terms of algebra, the reachability problem is equivalent to the subsemigroup membership problem. Several variants and generalization (rational subset membership problem, knapsack problem) have been recently investigated as well. Igor Potapov gave a survey talk on recent progress on the matrix reachability problem from a computer science perspective. Georg Zetzsche presented several new decidability results for the rational subset membership problem in wreath products. Moses Ganardi talked on wreath products as well, but put the focus on the knapsack problem.

The above talks and the open problem session identified several interersting open problems related to reachability problems. An outstanding open problem in this context asks whether the submonoid membership problem for the group $\mathrm{GL}_3(\mathbb{Z})$ is decidable. Recently it was shown that the submonoid membership problem for the Heisenberg group (a subgroup of $\mathrm{GL}_3(\mathbb{Z})$) is decidable, This result suggests two generalizations: (i) the rational subset membership problem for Heisenberg groups and (ii) the submonoid membership problem for groups of unitriangular integer matrices. In both case it is open whether the problem is decidable. Georg Zetzsche mentioned in his talk the submonoid membership problem and the rational subset membership problem in the Baumslag-Solitar group $\mathrm{BS}(1,2)$ as open research problems.

Concluding remarks and future plans. The seminar was well received as witnessed by the high rate of accepted invitations. There was a good balance between participants from computer science and pure mathematics, and this mixture led to many active discussions and the discovery of new connections and promising open problems. The organizers regard this seminar as a great success. With steadily increasing interactions between such researchers, we foresee another seminar focusing on the interplay between computer science and group theory. Finally, the organizers wish to express their gratitude to the Scientific Directors of the Dagstuhl Centre for their support of the seminar.

6.21 Users and automated driving systems: How will we interact with tomorrow's vehicles?

Organizers: Susanne Boll, Andrew L. Kun, Andreas Riener, and C. Y. David Yang

Seminar No. 19132

Date: March 24–29, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.111

Participants: Ignacio J. Alvarez, Martin Baumann, Susanne Boll, Linda Ng Boyle, Stephen Brewster, Duncan Brumby, Gary Burnett, Lewis Chuang, Birsen Donmez, Frank Flemisch, Joanne Harbluk, Christian P. Janssen, Wendy Ju, Josef Krems, Andrew Kun, Sabine Langlois, Roderick McCall, Alexander Meschtscherjakov, Ingrid Pettersson, Bastian Pfleging, Andreas Riener, Maria Rimini-Doering, Nele Rußwinkel, Albrecht Schmidt, Ronald Schroeter, Steven E. Shladover, Jaka Sodnik, Christine Sutter, Bruce N. Walker, C. Y. David Yang, Jürgen Ziegler

For much of the time since the invention of the automobile, human-machine interaction (HMI) in vehicles was reasonably clear: drivers controlled the vehicle by manipulating the steering wheel, pedals, and a few levers, buttons, or similar mechanical input devices [2]. They received information about the state of the vehicle through dials, warning lights, and sounds. And, they interacted with a relatively simple in-vehicle entertainment device: the radio, or perhaps the cassette- or CD player.

It is true that the number of input and output devices increased dramatically over the years — for example in the late 1950s, the Ford Edsel was described as a "devilish assemblage of gadgets" [5]. The Edsel was soon out of production, but the number of gadgets kept climbing. It is also true that drivers sometimes operated the vehicle when they were tired, and fell asleep at the wheel. Other times they consumed too much alcohol, and were not able to safely control their vehicle. Yet, the basic concepts of human-machine interactions in the vehicle were well-defined for research and development purposes. The driver's primary task was to drive: keep the vehicle on the road, avoid crashes, maneuver through traffic, and ultimately reach a destination. The driver also engaged in secondary tasks, such as manipulating the radio, as well as other non-driving-related tasks, such as talking to passengers, and eating. Creating good human-machine interfaces meant supporting the driver in these primary and secondary tasks, while assuring safety for everyone on the road.

Then, with the introduction of mobile computing devices, engagement in secondary tasks, such as talking to remote conversants, as well as sending text messages, and manipulating the interfaces of various mobile applications, became a significant issue in cars. In a sense these distractions were the same as those that drivers faced with the myriad buttons in the Ford Edsel. But, there were differences too: the Edsel did not allow the driver to communicate to remote conversants, nor did it have a touch-screen with ever-changing content.

Today, we again find ourselves at a crossroads. Our cars have myriad buttons, as well as different mobile technologies, both for drivers and for passengers. But, additionally, the primary task of driving is often shared between the driver and the vehicle [9]. Most studies in distracted driving tend to focus on how non-driving activities serve as a distraction from the primary task of vehicle control. In the context of highly automated vehicles (HAV), driving will be the distraction from non-driving activities [6]. Sometimes, the vehicle can effectively take over the driving task, and we can expect that this will become the norm rather than the exception in the foreseeable future: the driving task will be automated, at least on some road segments, and the driver will become a passenger. Yet, in this same foreseeable future we can also expect that the vehicle will sometimes hand the driving task back to the driver, who will have to transition back from the role of passenger to the role of the driver [14], [18]. This is the new landscape of in-vehicle human-machine interaction, and it presents a number of research questions that we addressed in this Dagstuhl seminar. In the rest of this report, we introduce pre-workshop tasks and summarize the activities and outcome of the seminar.

Automated traffic is a challenge not limited to the interaction between a human driver and an automated vehicle. Automated vehicles will be part of a mixed traffic with other traffic participants of less or no automation. Also further traffic participants such as pedestrian and bicycles are part of this and requires a certain level of communication and recognition of the vehicles intention and actions among vehicles and the surrounding traffic participants.

Research questions tackled in Dagstuhl seminar 19132

1. **Handover:** One of the key questions in designing in-vehicle human-machine interaction for partially automated vehicles is, how can the vehicle safely hand back the task of controlling the vehicle to the driver. In the short term this is one of the most important questions for those designing vehicles with automation, because in the short term such vehicles will have to hand control back to the driver relatively often [14], [15]. We need to understand how the modality, conveyed information, and reliability of take-over requests (TORs), engagement in non-driving tasks, and motion perception can influence drivers performance in task switching in highly automated driving context [3].
2. **Trust:** Drivers must trust the automation features in order to take advantage of them [19]. We need to individually understand the trust in the individual actions of the vehicle starting out from assistance systems [21] to more automated functions [13], [20].
3. **Creating a place for work and play:** One important benefit of automation would be that drivers can become passengers, and thus use the time in the vehicle to either be productive (work), or relax (play). How can human-machine interaction for automated vehicles be designed, such that drivers can take advantage of their newfound freedom from driving [9], [12]? How can we do this, taking into account the physical and computational characteristics of the vehicle, as well as the potential for motion sickness?
4. **Communication between all traffic participants:** With the advent of automation, the transportation environment will include partially and fully automated vehicles. Yet, manually driven vehicles will remain for the foreseeable future, as will pedestrians, bicyclists, and other transportation users. For safe driving, all of these transportation users will have to communicate, but it is not yet clear how this can best be accomplished [16].
5. **Advanced technologies for in-vehicle HMI:** The technologies that are available for human-machine interaction are continuously improving. Two exciting technologies that will be worth examining in the context of automated vehicles are speech interaction (e. g. [8]), and augmented reality e. g. [11] and [10].
6. **Legal aspects of in-vehicle interfaces:** Automation, as well as the user interfaces built for partially and fully automated vehicles, will have to fit into the legal structures of the countries where the vehicles are used [7]. What are these structures, what do designers need to know about them, and how can they help develop the future legal structures?

References

1 Ingrid Pettersson, Florian Lachner, Anna-Katharina Frison, Andreas Riener, and Andreas Butz. A bermuda triangle?: A review of method application and triangulation in user experience evaluation. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI '18, pages 461:1–461:16, New York, NY, USA, 2018. ACM.

2 Akamatsu, M., Green, P., & Bengler, K. Automotive technology and human factors research: Past, present, and future. In *International journal of vehicular technology*, 2013.

3 Borojeni, S.S., Weber, L., Heuten, W., & Boll, S. From reading to driving: priming mobile users for take-over situations in highly automated driving. *Paper presented at the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services*, 2018.

4 Borojeni, S. S., Chuang, L., Heuten, W., & Boll, S. Assisting Drivers with Ambient Take-Over Requests in Highly Automated Driving. *Paper presented at the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 2016.

5 Brooks, J. Business adventures: Twelve classic tales from the world of wall street: Open Road Media. 2014.

6 Hancock, P. A. Driven to distraction and back again. In *In J. Lee (Ed.), Driver distraction and inattention: Advances in research and countermeasures*, (pp. 9-25): CRC Press., 2013.

7 Inners, M., & Kun, A. L. Beyond Liability: Legal Issues of Human-Machine Interaction for Automated Vehicles. In *9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, (pp. 245-253), Oldenburg, Germany, 2017. ACM.

8 Kun, A., Paek, T., & Medenica, Z. The effect of speech interface accuracy on driving performance. *Paper presented at the Eighth Annual Conference of the International*, 2007.

9 Kun, A. L., Boll, S., & Schmidt, A. Shifting Gears: User Interfaces in the Age of Autonomous Driving. In *IEEE Pervasive Computing*, 15(1), 32-38., 2016.

10 Kun, A. L., Tscheligi, M., Riener, A., & Meulen, H. v. d. ARV 2017: Workshop on Augmented Reality for Intelligent Vehicles. *Paper presented at the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications Adjunct*, Oldenburg, Germany, 2017. ACM.

11 Kun, A. L., van der Meulen, H., & Janssen, C. P. Calling while driving: An initial experiment with HoloLens. *Paper presented at the Driving Assessment*, 2017.

12 Kun, A. L., Wachtel, J., Miller, W. T., Son, P., & Lavallière, M. ARV 2017: Workshop on Augmented Reality for Intelligent Vehicles. *Paper presented at the 7th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, 2015.

13 Löcken, A., Heuten, W., & Boll, S. AutoAmbiCar: Using Ambient Light to Inform Drivers About Intentions of Their Automated Cars. *Paper presented at the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications Adjunct* 2016.

14 Merat, N., Jamson, A. H., Lai, F. C., Daly, M., & Carsten, O. M. User interfaces for first responder vehicles: views from practitioners, industry, and academia. In *Transportation research part F: traffic psychology and behaviour*, 27, 274-282., 2014.

15 Mok, B., Johns, M., Miller, D., & Ju, W. Tunneled In: Drivers with Active Secondary Tasks Need More Time to Transition from Automation. *Paper presented at the 2017 CHI Conference on Human Factors in Computing Systems*, CHI , 2017.

16 Rasouli, A., Kotseruba, I., & Tsotsos, J. K. Agreeing to Cross: How Drivers and Pedestrians Communicate. In *arXiv preprint*, arXiv:1702.03555., 2017.

17 Riener, A., Boll, S., & Kun, A. L. Automotive User Interfaces in the Age of Automation (Dagstuhl Seminar 16262). *Paper presented at the Dagstuhl Reports.*, 2016.

18 van der Meulen, H., Kun, A. L., & Janssen, C. P. Switching Back to Manual Driving: How Does it Compare to Simply Driving Away After Parking? *Paper presented at the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, Ann Arbor, MI., 2016.

19 Wintersberger, P., Riener, A. Trust in Technology as a Safety Aspect in Highly Automated Driving. In *i-com* 15(3), 2016.

20 Wintersberger, P., Frison, A.-K., Riener, A., & Boyle, L. N. Towards a Personalized Trust Model for Highly Automated Driving. *Paper presented at the Mensch und Computer 2016–Workshopband*, 2016.

21 Yan, F., Eilers, M., Lüdtke, A., & Baumann, M. Developing a model of driver's uncertainty in lane change situations for trustworthy lane change decision aid systems. *Paper presented at the Intelligent Vehicles Symposium* (IV), IEEE, 2016.

L

(To the tune of L-O-V-E by Nat King Cole, Lyrics by Matthew Katzman)

G Em7 Am7 D7
L is for the likes of graph theory
Am7 D7 G
L is known to be contained in P
G7 C
L is very, very extraordinary
A A7 D A7 D7
L can even compute (u)s-t connectivity and

G Em7 Am7 D7
Logs are all that L will give to you
Am7 D7 G
Logs, but add a random bit or two
G7 C
Then RL we make it Hope a PRG will break it
G D7 (Am7) G
Logs were made for me and you

A F#m7 Bm7 E7
L has low circuit complexity
Bm7 E7 A
L is known to be in NC3 (and 2)
A7 D
L is very, very revolutionary
B B7 E B7 E7
L might even possibly be equal to NP and

A F#m7 Bm7 E7
Logs are all that L will give to you
Bm7 E7 A
Logs, the input length cannot accrue
A7 D
So we can't yet use ad- vice to complement RL: sad!
A E7 (Bm7) A
Logs were made for me and you!

Fig. 6.2
Song text written during Dagstuhl Seminar 19121 – Computational Complexity of Discrete Problems.

6.22 Programmable Network Data Planes

Organizers: Gianni Antichi, Theophilus Benson, Nate Foster, Fernando M. V. Ramos, and Justine Sherry

Seminar No. 19141

Date: March 31–April 5, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.3.178

Participants: Gianni Antichi, Mario Baldi, Sujata Banerjee, Theophilus Benson, Roberto Bifulco, Gordon Brebner, Marco Chiesa, Paolo Costa, Jon Crowcroft, Lars Eggert, Anja Feldmann, Andy Fingerhut, Nate Foster, Soudeh Ghorbani, Timothy G. Griffin, Stephen Ibanez, Changhoon Kim, Daehyeok Kim, Eder L. Fernandes, Alberto Lerner, Dotan Levi, Nick McKeown, Aurojit Panda, Justin Pettit, Ben Pfaff, Salvatore Pontarelli, Costin Raiciu, Fernando M. V. Ramos, Gabor Retvari, Hugo Sadok, Justine Sherry, Salvatore Signorello, Alexandra Silva, Robert Soulé, Alex Sprintson, David L. Tennenhouse, Laurent Vanbever, Stefano Vissicchio, David Walker, Hakim Weatherspoon, Minlan Yu, Noa Zilberman

Traditional networks are complex and hard to manage. It is difficult to configure networks according to predefined policies, and to reconfigure them in response to dynamic changes. Traditional networks are also vertically integrated: the control and data planes are bundled together. Around 10 years ago, the Software-Defined Networking (SDN) paradigm emerged and began to change this state of affairs. SDN breaks the vertical integration, separating the network's control logic from the underlying routers and switches (by means of a protocol such as OpenFlow) and promoting the (logical) centralization of network control. As such, it enabled the introduction of new abstractions in networking giving the ability to program the control plane of networks. Modern data center networks employ SDN-based techniques to simplify network management and operate at very large scale, and new networking services are now made possible – prominent examples are VMware's Network Virtualization Platform, Google's Andromeda and Microsoft's AccelNet.

Despite offering programmatic control to network operators, the original SDN data plane was limited to the protocols supported by OpenFlow. Over time, the OpenFlow specification evolved to support operator requirements, growing from 12 header fields in the original version to nearly 50 protocols in recent versions. The primary reason that OpenFlow is limited to specific "baked in" protocols is that the capabilities of switching chips are fixed at fabrication time. However, recent chip designs have demonstrated that it is possible to increase the flexibility of switch ASICs while still maintaining the terabit speeds required of networking hardware. In addition, as programming these chips is difficult – they expose their own low-level interface, akin to microcode programming – a domain-specific language, P4, was recently proposed to program network data planes (see p4.org). These advances are leading to a growing understanding of the inherent challenges related to data plane programming, resulting in further changes that promote future advances. For example, P4 was originally based on a simple architectural model, but has evolved to allow different switch architectures, aiming for stability of the language while increasing the flexibility to switch designers.

At the same time as programmable switches and programming languages such as P4 were being developed, a different group of researchers within the networking community has explored an alternative approach in which advanced data plane functionality is implemented on end hosts. This approach is often known as Network Function Virtualization (NFV). Platforms such as OpenVSwitch and Intel's DPDK framework make it possible to implement sophisticated packet-processing functions on end hosts rather than network switches, at line rates up to 10Gb/s and beyond. A key advantage of using CPUs is their flexibility, which makes it easy to adapt as requirements evolve. For example, it is straightforward to implement fine-grained monitoring of network flows or cryptographic operations – two pieces of functionality that are difficult to implement on standard switch ASICs.

In this context, the seminar on programmable data planes brought together leading practitioners from the areas of networking, systems, programming languages, verification, and hardware, to exchange ideas about important problems and possible solutions, and to begin the task of developing a research agenda related to programmable data planes. We have discussed several topics, including data plane architectures; programming languages, compilers and targets; use cases and applications; verification tools and formal methods; and end-system issues.

In the seminar we discussed questions including where packet-processing functionality should reside, how programmable data planes should evolve, how networks can benefit from these new elements, and how they can cope with the new challenges that arise. The focus was on the key challenges of the field and on the most fundamental problems to look at in the next 10 years, together with an aim to identify the "right" steps to take to move forward and the key problems to tackle next.

We have made some progress toward answering the following synergistic research questions during the seminar: What is the right division of labor between control and data plane? What are the right high-level language abstractions for programming networks, and what uarantees could we expect a compiler to provide reachability, security, or even properties as detailed as the correct use of cryptography? What is the trade-off between making more intelligent data plane architectures and the resulting network performances? Can we enhance current methods adopted to check network configuration errors with new solutions that automatically assure the absence of misconfiguration?

In the rest of the report we summarise the outcome of the most relevant discussions we had during the seminar.

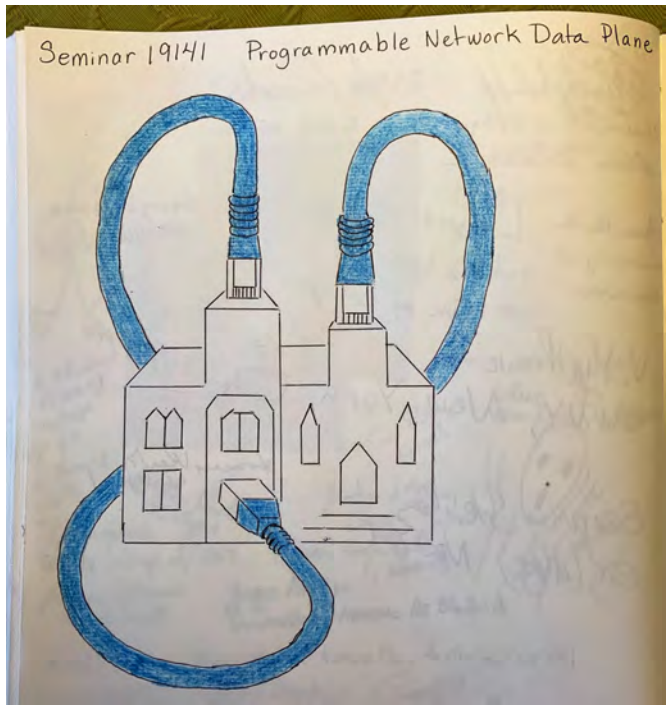

Fig. 6.3
"Tanya Fingerhut's awesome logo for our @dagstuhl seminar" Twitter post by 19141 Dagstuhl Seminar participant Nate Foster. https://twitter.com/natefoster/status/1113552318987939841. Photo courtesy of Nate Foster.

6.23 Visual Computing in Materials Sciences

Organizers: Christoph Heinzl, Robert Michael Kirby, Stepan V. Lomov, Guillermo Requena, and Rüdiger Westermann

Seminar No. 19151

Date: April 7–12, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.4.1

Participants: Amal Aboulhassan, Jan De Beenhouwer, Francesco De Carlo, Thomas Ertl, Christian Gollwitzer, Eduard Gröller, Doga Gursoy, Hans Hagen, Marcus Hanwell, Ulf Hassler, Hans-Christian Hege, Wolfgang Heidrich, Christoph Heinzl, Johann Kastner, Robert Michael Kirby, Fabien Leonard, Stepan V. Lomov, Lucia Mancini, Torsten Möller, Rajmund Mokso, Klaus Mueller, Vijay Natarajan, Ahmad Osman, Sidnei Paciornik, Kristi Potter, Bernhard Preim, Guillermo Requena, Gerik Scheuermann, Katja Schladitz, Johanna Schmidt, Jeff Simmons, Federico Sket, Daniela Ushizima, Daniel Weiskopf, Rephael Wenger, Rüdiger Westermann, Thomas Wischgoll

In this Dagstuhl workshop, we brought together computer and computational scientists interested in building tools for use in visual computing with material scientists with expressed interest in using such tools. As would be anticipated when one brings together two distinct fields, the initial challenge we encountered was that of language. Although both groups came together having experiences with visual computing tools – some as developers and some as users – although they often used the same terms, they semantically meant different things. We found that the Dagstuhl philosophy of "immersion" was most helpful to this issue as having several days together helped break down these barriers. Over the course of the week, we interspersed talks by computational scientists and material scientists. The talks by computational scientists often presented their current understanding of what kinds of tools are needed, demonstrations of current tools they have developed in collaboration with domain-specific experts, and success stories of applications they have currently impacted. The talks by the material scientists often presented a description of the tools they currently use, the positive points and deficiencies of current tools, the types of features that they would like to see in future tools, and examples of current challenge problems and how they might be impacted by the next generation of tools.

Fundamental Results:

1. The systems that are desired by many material scientists will be used both for exploration and for interactive steering. When used for exploration, material scientists want tools that not only present the data with its corresponding reliability (uncertainty) bounds, but which also give predictive capabilities such as where next to sample.
2. There is a general acknowledgement that both automation and interactivity are needed. Automation of tasks and procedures through AI and Machine Learning can be used to help deal with the volumes of data being produced – helping scientists sift through the field of possibilities to isolate those places for which they should expend human effort. At the same time, there are many current practices that continue to require "the human in the loop" to make decisions. In such cases, tools are needed that have smart defaults but yet allow the user to explore, navigate and possibly refine data.
3. Although many current tools used for material science applications leverage previous visualization and interaction technologies, there is still much to be done. Many material science applications require specialization of currently existing algorithms and techniques, especially in cases of real-time systems. Furthermore, many techniques originally designed for batch or manual processing need to be re-engineered to allow for the interactive procedures required by current and future material science application scientists.
4. With regards to visualization scientists, there is a need for both data and tasks. Many researchers requested data on which they can try their methods. In addition to the data itself, descriptors of the data are necessary so that it can be interpreted properly. Once read into their system, the visualization scientists then requested a collection of tasks (driven by the material science domain experts) which would help drive their tool development and evaluation.

Final Comments Due to the ever-increasing interest in this topic, we foresee that future review articles and/or special issues of journals driven by multilateral research cooperations between seminars' participants will be an outcome of this workshop. To ensure and stimulate further cooperation in this field, a list of specific follow up activities has been elaborated and discussed with the participants. All in all, a fruitful discussion was stimulated across the two domains throughout the complete week of this Dagstuhl workshop which will become more obvious in joint research efforts of all kinds.

6.24 Emerging Hardware Techniques and EDA Methodologies for Neuromorphic Computing

Organizers: Krishnendu Chakrabarty, Tsung-Yi Ho, Hai Li, and Ulf Schlichtmann

Seminar No. 19152

Date: April 7–10, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.4.43

Participants: Krishnendu Chakrabarty, Meng-Fan Chang, Jian-Jia Chen, Yiran Chen, Federico Corradi, Rolf Drechsler, Deliang Fan, Tsung-Yi Ho, Alex Pappachen James, Bing Li, Hai Li, Darsen Lu, Christoph Maier, Onur Mutlu, Qinru Qiu, Garrett S. Rose, Yulia Sandamirskaya, Johannes Schemmel, Ulf Schlichtmann, Yu Wang, Chia-Lin Yang

The explosion of *big data* applications imposes severe challenges of data processing speed and scalability on traditional computer systems. However, the performance of von Neumann architecture is greatly hindered by the increasing performance gap between CPU and memory, motivating active research on new or alternative computing architectures. Neuromorphic computing systems, that refer to the computing architecture inspired by the working mechanism of human brains, have gained considerable attention. The human neocortex system naturally possesses a massively parallel architecture with closely coupled memory and computing as well as unique analog domain operations. By imitating this structure, neuromorphic computing systems are anticipated to be superior to conventional computer systems across various application areas. In the past few years, extensive research studies have been performed on developing large-scale neuromorphic systems. Examples include IBM's TrueNorth chip, the SpiNNaker machine of the EU Human Brain Project, the BrainScaleS neuromorphic system developed at the University of Heidelberg, Intel's Loihi etc. These attempts still fall short of our expectation on energy-efficient neuromorphic computing systems with online, real-time learning and inference capability. The bottlenecks of computation requirements, memory latency, and communication overhead continue to be showstoppers. Moreover, there is a lack of support in design automation of neuromorphic systems, including functionality verification, robustness evaluation and chip testing and debugging. Hardware innovation and electronic design automation (EDA) tools are required to enable energy-efficient and reliable hardware implementation for machine intelligence on cloud servers for extremely high performance as well as edge devices with severe power and area constraints.

The goal of the seminar was to bring together experts from different areas in order to present and to develop new ideas and concepts for emerging hardware techniques and EDA methodologies for neuromorphic computing. Topics that were discussed included:

- Neuroscience basics
- Physical fundamentals
- New devices and device modeling
- Circuit design and logic synthesis
- Architectural innovations
- Neurosynaptic processor and system integration
- Design automation techniques
- Simulation and emulation of neuromorphic systems
- Reliability and robustness
- Efficiency and scalability
- Hardware/software co-design
- Applications

The seminar facilitated greater interdisciplinary interactions between physicists, chip designers, architects, system engineers, and computer scientists. High-quality presentations and lively discussions were ensured by inviting carefully selected experts who participated in the seminar. All of them have established stellar reputations in the respective domains. As a result, we developed a better understanding of the respective areas, generated impetus for new research directions, and ideas for areas that will heavily influence research in the domain of neuromorphic design over the next years.

At the end of the seminar, we identified the following four areas as being among the most important topics for future research: *computing-in-memory*, *brain-inspired design and architecture*, *new technologies and devices*, and *reliability and robustness*. These research topics are certainly not restricted to and cannot be solved within one single domain. It is therefore imperative to foster interactions and collaborations across different areas.

6

6.25 Ethics and Trust: Principles, Verification and Validation

Organizers: Michael Fisher, Christian List, Marija Slavkovik, and Astrid Weiss

Seminar No. 19171

Date: April 22–26, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.4.59

Participants: Andrea Aler Tubella, Jan M. Broersen, Einar Duenger Bøhn, Raja Chatila, Emily Collins, Louise A. Dennis, Franz Dietrich, Clare Dixon, Hein Duijf, Abeer Dyoub, Sjur K. Dyrkolbotn, Kerstin I. Eder, Kerstin Fischer, Michael Fisher, Marc Hanheide, Holger Hermanns, John F. Horty, Maximilian Köhl, Robert Lieck, Felix Lindner, Christian List, Bertram F. Malle, Andreas Matthias, AJung Moon, Marcus Pivato, Thomas Michael Powers, Teresa Scantamburlo, Munindar P. Singh, Marija Slavkovik, Kai Spiekermann, Myrthe Tielman, Suzanne Tolmeijer, Leon van der Torre, Astrid Weiss, James E. Young

Academics, engineers, and the public at large, are all wary of *autonomous systems*, particularly robots, drones, "driver-less" cars, etc. Robots will share our physical space, and so how will this change us? With the predictions of roboticists in hand, we can paint portraits of how these technical advances will lead to new experiences and how these experiences may change the ways we function in society. Two key issues are dominant once robot technologies have advanced further and yielded new ways in which we and robots share the world: (1) will robots behave *ethically*, i.e. as we would want them to, and (2) can we *trust* them to act to our benefit. It is more these barriers concerning ethics and trust than any engineering issues that are holding back the widespread development and use of autonomous systems. One of the hardest challenges in robotics is to reliably determine desirable and undesirable behaviours for robots. We are currently undergoing another technology-led transformation in our society driven by the outsourcing of decisions to intelligent, and increasingly autonomous, systems. These systems may be software or embodied units that share our environment. The decisions they make have a direct impact on our lives. With this power to make decisions comes the responsibility for the impact of these decisions – legal, ethical and personal. But how can we ensure that these artificial decision-makers can be *trusted* to make safe and *ethical* decisions, especially as the responsibility placed on them increases?

The related previous Dagstuhl Seminar 16222 on *Engineering Moral agents: From human morality to artificial morality* in 2016, highlighted further important areas to be explored, specifically:

- the extension of 'ethics' to also address issues of 'trust';
- the practical problems of implementing ethical and trustworthy autonomous machines;
- the new verification and validation techniques that will be required to assess these dimensions.

Thus, we thought that the area would benefit from a follow-up seminar which broadens up the scope to Human-Robot Interaction (HRI) and (social) robotics research.

We conducted a four-day seminar (1 day shorter than usual due to Easter) with 35 participants with diverse academic backgrounds including AI, philosophy, social epistemology, Human-Robot Interaction, (social) robotics, logic, linguistics, political science, and computer science. The first day of the seminar was dedicated to seven invited 20-minute talks which served as tutorials. Given the highly interdisciplinary nature of the seminar, the participants from one discipline needed to be quickly brought up to speed with the state of the art in the discipline not their own. Moreover, the goal of these tutorials was to help develop a common language among researchers in the seminar. After these tutorials we gave all participants the chance to introduce their seminar-related research in 5-minute contributed talks. These talks served as a concise way to present oneself and introduce topics for discussion.

Based on these inputs four topics were derived and further explored in working groups through the rest of the seminar: (1) Change of trust, including challenges and methods to foster and repair trust; (2) Towards artificial moral agency; (3) How do we build practical systems involving ethics and trust? (2 sub-groups) (4) The broader context of trust in HRI: Discrepancy between expectations and capabilities of autonomous machines. This report summarizes some of the highlights of those discussions and includes abstracts of the tutorials and some of the contributed talks. Ethical and trustworthy autonomous systems are a topic that will continue to be important in the coming years. We consider it essential to continue these cross-disciplinary efforts, above all as the seminar revealed that the "interactional perspective" of the "human-in-the-loop" is so far underrepresented in the discussions and that also broadening the scope to STS (Science and Technology Studies) and sociology of technology scholars would be relevant.

6.26 Computational Creativity Meets Digital Literary Studies

Organizers: Tarek Richard Besold, Pablo Gervás, Evelyn Gius, and Sarah Schulz

Seminar No. 19172

Date: April 22–25, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.4.87

Participants: Leonid Berov, Tarek Richard Besold, Amilcar Cardoso, João Miguel Cunha, Thierry Declerck, Mark Finlayson, Pablo Gervás, Evelyn Gius, Marina Grishakova, Christian Guckelsberger, Christopher Hench, Kai-Uwe Kühnberger, Jonas Kuhn, Oliver Kutz, Carlos León, Rafael Pérez y Pérez, Enric Plaza, Manuel Portela, Dino Pozder, Nils Reiter, Sarah Schulz, Hannu Toivonen, Sara L. Uckelman, Tony Veale, Philipp Wicke, Sina Zarrieß

Literary studies (LS) is a subfield of the humanities that provides a diversity of possible views on its objects of investigation. The universal approach to literary texts does not exist, instead there are many, sometimes incompatible theories that can be applied for the interpretation of literary texts. Additionally, with the emerging of the Digital Humanities (DH) the deployment of computational methods has been introduced into LS, leading to a further expansion of the range of theories and methodologies of text analysis and interpretation. Against that backdrop in the last decade much effort has especially been put into developing approaches that cover rather complex concepts for text analysis, including, among other, network theory (e.g., [7]) and approaches from distributional semantics for topic modelling and word vector estimation (e.g., [8]). This considerably changed the prerequisites of DH research in the field of LS. In many cases it is no longer possible to simply apply a predefined tool or algorithm, requiring traditionally trained LS scholars to move away from their disciplinary paradigm of individual research and towards adapting collaborative modes that can provide both LS and computational expertise. Researchers in Natural language processing (NLP) have shown considerable interest in text-based DH research. This interest is not only motivated by the diversity and complexity of the research questions, which offers an ideal testbed for the development of new methods and combined workflows, but also by the nature of texts found in the context of these research questions which are often diverse with respect to their lexical and syntactic range – meeting the need for this type of data in work aiming for more flexible NLP approaches. Computational Creativity (CC) is a multidisciplinary endeavour, modelling, simulating or replicating aspects of creativity using a computer, in order to achieve one of several ends: Either to construct a program or computer capable of human-level creativity, or to better understand human creativity and to formulate an algorithmic perspective on creative behaviour in humans, or to design programs that can enhance human creativity without necessarily being creative themselves (a concise overview of the main aspects of the field has, for instance, been laid out by [1]). One of CC's most popular subfields is Computational Storytelling (CS), where researchers hitherto have mainly thought about the structure and logical implications of building blocks of stories, leaving most other dimensions of narrative construction out of consideration.

Taking stock of this overall state of affairs and the specific situation in the respective fields, the seminar was constructed around several main challenges:

- One of the major challenges in DLS is the approximation of concepts with computational approaches to, i.e. their operationalization, that not only requires a translation of the concepts, but also a deep understanding of the deployed computational approaches used. This gap can be tackled best by providing expertise from the fields concerned. Whereas NLP is already accepted as such a field (but still needed), CS has not been taken much into consideration yet. A second type of collaboration that still needs to be intensified is the one that connects the interpretative, manual annotations from DLS (e.g., [4]) with computational approaches to text analysis and generation.
- NLP has focused on a limited variety of texts in its beginnings and suffers from a bias towards newspaper texts. Even though there are efforts towards more diverse and flexible text processing, the constant lack of data is a problem. DH and CC offer a variety of texts to improve this situation – but are hitherto underused in that capacity.

- CC, CS focuses almost exclusively on plot and logical structure of storytelling. However, a narrative is a complex web of different factors that are well-investigated in classical disciplines. While much work is based on formalist theories about narrative (especially [9]), other approaches from narrative theory still need to be explored better. For example, CS could benefit from the well-established fields of semiotics (e.g., [5]) and structuralism (e.g., [3]) as well as from more recent, reader-oriented developments in cognitive and empirical narratology (e.g. [6]; [2]).

In order to make researchers from the participating communities a) aware of the challenges and the corresponding opportunities an interdisciplinary meeting like the seminar offered, and b) make them take advantage of these opportunities still on-site, the seminar was split between presentations from researchers describing their recent work and questions they wanted to highlight for the audience, and "hackathon" phases in which decidedly interdisciplinary teams of participants worked on concrete projects.

The following pages summarize the content of these presentations and the outcomes of the group projects.

References

1 M. A. Boden. How computational creativity began. In M. Besold, T. R.; Schorlemmer and A. Smaill, editors, *Computational Creativity Research: Towards Creative Machines*. Atlantis Press, Amsterdam, 2015.

2 M. Bortolussi, P. Dixon, and F. C. E. P. Dixon. *Psychonarratology: Foundations for the Empirical Study of Literary Response*. Psychonarratology: Foundations for the Empirical Study of Literary Response. Cambridge University Press, 2003.

3 G. Genette. *Narrative discourse:*. G – Reference, Information and Interdisciplinary Subjects Series. Cornell University Press, 1980.

4 E. Gius. In *Diegesis*, page 4. 2016.

5 A. J. Greimas. *Structural Semantics: An Attempt at a Method*. University of Nebraska Press, 1983.

6 D. Herman. *Story Logic: Problems and Possibilities of Narrative*. Frontiers of narrative. University of Nebraska Press, 2002.

7 F. Moretti. *Network Theory, Plot Analysis*. Literary lab. Stanford Literary Lab, 2011.

8 Christof Schöch. Topic Modeling Genre: An Exploration of French Classical and Enlightenment Drama. *Digital Humanities Quarterly*, November 2016. This is a pre-publication version of an article to appear in Digital Humanities Quarterly. Last revision: October 2016.

9 Vladimir Propp. *Morphology of the Folktale*. University of Texas Press, 2010.

6.27 Computational Geometry

Organizers: Siu-Wing Cheng, Anne Driemel, and Jeff Erickson

Seminar No. 19181

Date: April 28–May 3, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.4.107

Participants: Mikkel Abrahamsen, Peyman Afshani, Kevin Buchin, Maike Buchin, Hsien-Chih Chang, Siu-Wing Cheng, Man-Kwun Chiu, Arnaud de Mesmay, Anne Driemel, Alperen Ergür, Jeff Erickson, Esther Ezra, Kyle Jordan Fox, Marc Glisse, Rolf Klein, Stefan Langerman, Maarten Löffler, Kurt Mehlhorn, Wouter Meulemans, Guillaume Moroz, David M. Mount, André Nusser, Eunjin Oh, Zuzana Patáková, Benjamin Raichel, Natan Rubin, Maria Saumell, Lena Schlipf, Raimund Seidel, Micha Sharir, Bettina Speckmann, Monique Teillaud, Hans Raj Tiwary, Marc van Kreveld, André van Renssen, Suresh Venkatasubramanian, Kevin Verbeek, Antoine Vigneron, Birgit Vogtenhuber, Nicola Wolpert

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 this 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 algebraic methods in computational geometry as well as geometric data structures. The former focus area has seen exciting recent progress and the latter is a fundamental topic at the heart of computational geometry. There are numerous opportunities for further cross-disciplinary impact.

Algebraic Methods in Computational Geometry. The polynomial method of Guth and Katz of 2010 has had a fundamental impact on discrete geometry and other areas, which was already envisioned by the talk of Jiří Matoušek at the Annual European Workshop on Computational Geometry in 2011, four years before he passed away. Indeed, the polynomial method has attracted the attention of many researchers, including famous ones like Janos Pach, Micha Sharir, and Terence Tao. Applications have been found not only in making progress on long-standing combinatorial geometry problems, but also in the design and analysis of efficient algorithms for fundamental geometric problems such as range searching, approximate nearest search, diameter, etc. The polynomial method is very powerful and it offers a new research direction in which many interesting new results can potentially be discovered.

Geometric Data Structures. Many beautiful results in geometric data structures have been established in the early days of the field. Despite of this, some long-standing problems remain unresolved and some of the recent progress is in fact made using the polynomial method mentioned previously. Independently, there have been some recent advances in our understanding of lower bounds and the usage of more sophisticated combinatorial constructions and techniques such as shallow cuttings, optimal partition trees, discrete Voronoi diagrams, etc. There are also new applications that require the modeling of uncertain data and hence call for a study of the performance of geometric data structures under a stochastic setting.

6

6.28 Multi-Document Information Consolidation

Organizers: Ido Dagan, Iryna Gurevych, Dan Roth, and Amanda Stent

Seminar No. 19182

Date: April 28–May 3, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.4.124

Participants: Omri Abend, Sebastian Arnold, Timothy Baldwin, Jonathan Berant, Giuseppe Carenini, Ido Dagan, Dipanjan Das, Daniel Deutsch, Laura Dietz, Yoav Goldberg, Dan Goldwasser, Iryna Gurevych, Heng Ji, Ayal Klein, Alexander Koller, Chin-Yew Lin, Fei Liu, Nafise Sadat Moosavi, Barbara Plank, Nils Reimers, Dan Roth, Steve S. Skiena, Gabriel Stanovsky, Amanda Stent, Ivan Titov, Kentaro Torisawa, Gisela Vallejo, Andreas Vlachos, Yue Zhang

Today's natural language processing (NLP) systems mainly work on individual text pieces like individual sentences, paragraphs, or documents. For example, most question answering systems require that the answer to a user's questions is provided in a single document, ideally in a single sentence. If the information is scattered across documents, most systems will fail. The capability of current systems to link information across multiple documents is often limited.

This is in strong contrast to how humans answer difficult questions or make complex decisions. We usually read multiple documents on a topic and then infer the answer to the question or we make a decision based on the evidence we found. In most cases, we consolidate the information across multiple sources. Further, considering only one document can create a biased or incomplete view on a topic. Many aspects in our life are open for multiple interpretations and each author must limit which and how to present information in a document. By reading multiple documents, we are able to identify overlaps, differences, and opposing views between authors. Considering and merging these possible opposing views can be a crucial step in everydays decision making. For example, when booking a hotel, one might read multiple user reviews and create an internal understanding of positive and negative aspects of the hotel.

At this 5-day Dagstuhl Seminar, an interdisciplinary collection of leading researchers discussed and develop research ideas that will lead to advanced multi-document information consolidation systems and enable modern NLP systems to profit from a multi-document perspective.

The seminar was centered around four major themes:

1. how to represent information in multi-document repositories;
2. how to support inference over multi-document repositories;
3. how to summarize and visualize multi-document repositories for decision support; and
4. how to do information validation on multi-document repositories.

Questions of semantics, pragmatics (author perspectives, argumentation), representation, and reasoning (including spatio-temporal reasoning and entailment) arose across these themes.

Information Representations and Inference are the theoretical foundation that allows systems to extract information from multiple documents and to infer new knowledge. The challenge is to find a representation that can broadly be used. Multiple documents are likely to bring up multiple perspectives and identifying the relations between them is at the heart of multi-document inference.

A connection to real applications, used in actual user scenarios, is critical for the advancement of the multi-document information consolidation field. Multi-document systems are especially useful in situations where users must make complex decisions. In such situations, users often search for sources that provide information or arguments for or against certain decisions. Hence, one working group focused on Multi-Document Systems in User Decision Scenarios. In order to provide value to users, the systems must return true statements (accurate syntheses) given all the available context. Otherwise, the user lose their trust in the system. However, the internet is full of statements that are intentionally or unintentionally misleading. So how do we identify these misleading statements and avoid that those are presented to a user without the necessary context? This research question was addressed by a working group focusing on Information Validation for Multi-Document Scenarios.

Seminar participants, including established experts and promising young researchers from academia and industry, had the opportunity to present research ideas, to outline their vision regarding the future of multi-document information consolidation

technologies, and to collaborate in discussion groups led by the seminar organizers.

Each seminar participant joined two themes with regular cross-theme meetings. As the topics are quite novel in the research community, no established terminology and task definition exists. Hence, participants discussed how these tasks can be defined such that these can be scientifically studied. For example, what does it mean to validate a claim? The participants discussed issues with existing approaches and proposed new research topics, that could be the content of a Ph.D. thesis.

The last day of the seminar was used to summarize results and to create collaborations for future research projects. In total, 12 joint research ideas were proposed. For most of the ideas, this is a new collaboration.

Fig. 6.4
"What better place to start my #sabbatical year than the oval office at wonderful @dagstuhl? I feel privileged. First goal: prepare a talk for @DAAD_Germany RISE meeting in Heidelberg (Research Internships in Science and Engineering) #amworkingonit" Twitter post by 19269 Dagstuhl research guest Christoph Becker. https://twitter.com/ChriBecker/status/1145707368967225345. Photo courtesy of Christoph Becker.

6.29 Software Evolution in Time and Space: Unifying Version and Variability Management

Organizers: Thorsten Berger, Marsha Chechik, Timo Kehrer, and Manuel Wimmer

Seminar No. 19191

Date: May 5–10, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.5.1

Participants: Sofia Ananieva, Sven Apel, Don Batory, Thorsten Berger, Danilo Beuche, Paulo Borba, Götz Botterweck, Andreas Burger, Marsha Chechik, Paul Grünbacher, Timo Kehrer, Heiko Klare, Anne Koziolek, Jacob Krüger, Ralf Lämmel, Yi Li, Lukas Linsbauer, Henrik Lönn, Wardah Mahmood, Shahar Maoz, Mukelabai Mukelabai, Sarah Nadi, Iris Reinhartz-Berger, Julia Rubin, Ina Schaefer, Klaus Schmid, Sandro Schulze, Christoph Seidl, Ramesh Sethu, Stefan Stanciulescu, Daniel Strüber, Gabriele Taentzer, Leopoldo Teixeira, Thomas Thüm, Mahsa Varshosaz, Eric Walkingshaw, Andrzej Wasowski, Bernhard Westfechtel, Manuel Wimmer, Shurui Zhou

Overview and Motivation

Modern software systems evolve rapidly and often need to exist in many variants. Consider the Linux kernel with its uncountable number of variants. Each variant addresses different requirements, such as runtime environments ranging from Android phones to large super-computers and server farms. At the same time, the Linux kernel frequently boasts new versions, managed by thousands of developers. Yet, software versions–resulting from evolution in time–and variants–resulting from evolution in space–are managed radically differently. Version management relies on a version control system (Git) and sophisticated workflows–concepts that have been developed for decades in the field of software configuration management (SCM) [13, 24, 25]. Variant management in the Linux kernel relies on techniques known from the field of software product line engineering (SPLE) [12, 14, 28], such as an integrated software platform, a variant-aware build system [8], an interactive configurator tool [31], and a model-based representation [2, 9, 10, 18] of all kernel features [5, 29]. The Linux kernel is exemplary for many large-scale, variant-rich, and rapidly evolving software systems in industry [4, 6, 33], especially in the domains of embedded, cyber-physical, automotive, and avionics control systems.

Despite decades of research in both fields, the effective evolution of variant-rich systems is still an open problem. Three main challenges exist. First, while version control systems are well-integrated into development processes, product-line engineering requires investment into additional tooling and different processes that are difficult to adopt. In fact, organizations rarely adopt product line engineering from scratch [7], but rather use readily available version control systems with their branching and forking facilities–a strategy known as clone & own [11, 16]. While this strategy is simple, it does not scale with the number of variants, and then requires evolving (i.e., re-engineering) cloned variants into a product-line platform [3]. Second, evolving product-line platforms is substantially more complex than evolving single variants, mainly since developers need to work on all variants at the same time [26]. Third, the granularity of tracking versions of variants is still unclear. While the whole platform can be versioned, ideally, versioning at the level of features should be supported.

In summary, SCM and SPLE are two widely established, yet actively researched software engineering disciplines offering a variety of concepts to deal with software versions and variants [15, 17, 19, 22]. Yet, despite various attempts [21, 23, 34, 35], none of the two disciplines has been successful in establishing unified solutions addressing both problems at the same time–mainly due to the isolation of both communities and due to the absence of realistic and widely accepted requirements on how to evaluate the effectiveness of techniques for managing both versions and variants.

Goals of the Seminar

This Dagstuhl Seminar aimed at establishing a body of knowledge on unified version and variant management. We invited leading practitioners and researchers from both disciplines to discuss each other's challenges, solutions, and experiences. The seminar's goals were to: (i) survey state-of-the-art SCM and SPLE concepts and map both areas' terminologies and open problems, (ii) gather industrial and academic challenges and requirements on integrated version and variant management, (iii) survey and assess existing evaluation approaches, and (iv) stablish a research agenda, research infrastructure, and working groups. To guide future research, the participants also discussed the basis to work on improved evaluation approaches–as benchmarks for new version and variant management techniques. As such, the long-term goal of the seminar was to enable the development and evaluation of enhanced version and variant management techniques that will be adopted in practice.

Week Overview

Monday. After an introduction of all participants, the seminar started off with general talks on versioning and variability. Bernhard Westfechtel set the stage with an introduction into version management concepts and workflows, which already illustrated some overlap with variability management concepts. For instance, directed deltas are conceptually similar to compositional variation mechanisms (e.g., feature modules or delta modules), and the construction of versions in intensional versioning can be related to the configuration-based derivation of individual variants from a product-line platform. The seminar continued with a talk by Don Batory, who discussed the integration of version control systems, variability management techniques, and integrated development environments (IDEs) based on ideas centering around a better representation and execution of program refactorings in versioned and variant-rich software systems. The talk by Thorsten Berger (actually given on Tuesday, since the introduction round and discussions for the other talks took more time) followed up on the concepts introduced in the previous talks and presented a survey on variation control systems, which support developers managing variant-rich systems in terms of features. Such variation control systems go back to the end of the 1970s with concepts and prototypes developed in the SCM community, but never made it into the mainstream. The talk surveyed their concepts and discussed problems likely prohibiting their adoption. Thereafter, we enjoyed three talks on industrial perspectives given by our industrial practitioners: Henrik Lönn (Volvo), Danilo Beuche (pure::systems), and Ramesh S. (General Motors; talk also given on Tuesday for timing reasons), confirming and explaining the gaps between academia and industry.

Tuesday. The day started with an introduction into the prospective breakout groups for the afternoon, followed by the talk of Christoph Seidl on versioning of product lines relying on a representation of feature versions in a new dialect of feature models, called Hyper Feature Models. Thereafter, the breakout sessions on four relevant topics took place, specifically: on a conceptual model to map SPLE and SCM concepts, on operations for managing versions and variants, on analyses of versions and variants, on workflows for managing versions and variants, and on first-class support of variability and versioning in programming languages. A benchmarking group was discussed, but abandoned in favor of first working on the foundations before discussing benchmarking techniques to evaluate prospective unified techniques for versioning and variability. The breakout group discussions continued until the afternoon, before the remaining talks from Monday were given (Thorsten Berger and Ramesh Sethu), followed by lightning talks from Shurui Zhou and Sandro Schulze. Shurui discussed the relevance of version and variability management in the domain of engineering AI-based systems, where models and large dataset need to be managed. Sandro proposed a round-trip-engineering process relying on unified management of versioning and variability, relying on automated extraction of variability information from cloned variants (which should be integrated into a platform in a round-trip-engineering manner).

Wednesday. We started the day with a talk by Daniel Strüber on benchmarking scenarios and a survey of existing benchmarks. In fact, it is a common consensus of the community that the lack of strong, landmark benchmarks hinders the progress in both communities (SCM and SPLE). Thereafter, Yi Li presented his work on slicing of the history of software codebases along features, where features are represented by test cases to help identifying the relevant code in a longitudinal manner. Thomas Thüm then presented a vision on the–ideally automated–synchronization of cloned variants as followed by the VariantSync research project which is led by Thomas and Timo Kehrer. Thomas also presented a very first prototypical implementation of the VariantSync tool. The approach shares, based on audience feedback, ideas with the Virtual Platform, proposed by researchers in 2014 [1]. In the afternoon, the majority of the participants continued their discussion on their group trip to the city of Trier and a dinner at a local winery.

Thursday. The day began with a talk by Gabriele Taentzer, presenting a generalizing framework for transformations of software product lines, relying on the formalism of category theory. Another talk was given by Julia Rubin on equivalence checking of variants based on behavior instead of structural characteristics of changes. Thereafter, the breakout groups continued their discussions until the later afternoon, where the results were presented to the other seminar participants. After dinner, two lightning talks were given by Paulo Borba and Iris Reinhartz-Berger. Paulo discussed the detection of semantic merge conflicts in the light of avoiding unwanted feature interactions, and Iris presented insights from two research projects on behavior-derived variability analysis and mechanisms recommendation.

Friday. The last day of the seminar started with a talk by Lukas Linsbauer on his work towards a feature-oriented and distributed version-control system, relying on the variant-integration tooling ECCO. We then had a closing discussion, re-iterating the main challenges we identified throughout the seminar, as well as discussing future work.

Outcome of the Seminar

The seminar established breakout groups who continued their discussion after the seminar and already published two papers [36, 37] at the VariVolution workshop, hosted at the Systems and Software Product Line Conference (SPLC). In addition, a paper accepted at the main track of SPLC on benchmarking, relying on input from the seminar participants via a survey [32], and providing an initial infrastructure for community-oriented benchmark creation,[50] can be seen as a core outcome of the seminar.

A core topic of the final discussion was the teaching of SPLE and SCM concepts–an important means to eventually improve the handling of versions and variants in practice. One of the problems identified is that, while SCM is covered sufficiently, the relevant variability-management concepts are not taught at the Bachelor's level in the majority of universities. However, the discussants believe that practicing feature-oriented analysis and design early in the curriculum would be beneficial, where currently object-oriented analysis and design is dominating. Interestingly, based on the experience of the discussants, SPLE is still seen as something rather futuristic by students, which is somewhat surprising, given that building highly configurable systems and software platforms are established practices, so perhaps there is a perception and awareness problem that teaching needs to address. Naturally, a course teaching SPLE at the Bachelor's level should also teach the relevant SCM concepts. A closely related topic discussed is that of teaching architectures, especially those of product lines, which is not really in the focus of current software architecture courses. Of course, it is generally difficult to talk to students about

[50] https://bitbucket.org/easelab/evobench

software architecture, since, as a discussant explains, a relevant abstract concept that students do not immediately perceive as relevant in the course of the studies. In contrast, with compilers and databases, students obtain some hands-on experience, which allows them to relate more closely to, especially with respect to a future job in the industry. This calls for close collaboration with industry in SPLE courses.

Establishing benchmarks turned out to be a more difficult problem than expected. Benchmarking was prominently discussed, as well as input elicited for a set of 11 high-level benchmarking scenarios defined by some of the seminar participants and organizers before the seminar. The participants plan to follow-up on creating concrete benchmarks upon the infrastructure created.[50] One idea is to build a web application to contribute specific benchmark data (e.g., code integration examples, comprising the original code variants and the final result as a ground truth) to establish a community benchmark. Another interesting comment was that the currently published case studies and experience reports about variability management and product lines are relatively old and do not provide sufficient technical details. Furthermore, they also do not highlight the problems associated with clone & own and the need for product-line migration techniques adequately. This discussion is a call to arms for improving the benchmarking situation in the SCM and SPLE community.

Last but not least, an important outcome of the final discussion session of the seminar is the need for a commonly agreed set of core concepts, mechanisms and practices–a well-documented Body of Knowledge (BOK) of our discipline. Currently, only some aspects of versioning in time and space are partially covered by the Software Engineering BOK (SWEBOK). However, for promoting a consistent view of our discipline worldwide and beyond our discipline borders as well as for having a foundation for a consistent curriculum development, a dedicated BOK or an extension of the SWEBOK may be necessary as a community effort.

References

1 Michal Antkiewicz, Wenbin Ji, Thorsten Berger, Krzysztof Czarnecki, Thomas Schmorleiz, Ralf Lämmel, Stefan Stanciulescu, Andrzej Wąsowski, and Ina Schaefer. Flexible product line engineering with a virtual platform. In *ICSE*, 2014.

2 Sven Apel, Don Batory, Christian Kästner, and Gunter Saake. *Feature-Oriented Software Product Lines*. Springer, Berlin Heidelberg, 2013.

3 Wesley K. G. Assunção, Roberto E. Lopez-Herrejon, Lukas Linsbauer, Silvia R. Vergilio, and Alexander Egyed. Reengineering legacy applications into software product lines: a systematic mapping. *Empirical Software Engineering*, 22(6):2972–3016, 2017.

4 Jonatas Ferreira Bastos, Paulo Anselmo da Mota Silveira Neto, Padraig O'Leary, Eduardo Santana de Almeida, and Silvio Romero de Lemos Meira. Software product lines adoption in small organizations. *Journal of Systems and Software*, 131(Supplement C):112–128, 2017.

5 Thorsten Berger, Daniela Lettner, Julia Rubin, Paul Grünbacher, Adeline Silva, Martin Becker, Marsha Chechik, and Krzysztof Czarnecki. What is a Feature? A Qualitative Study of Features in Industrial Software Product Lines. In *SPLC*, 2015.

6 Thorsten Berger, Divya Nair, Ralf Rublack, Joanne M. Atlee, Krzysztof Czarnecki, and Andrzej Wasowski. Three cases of feature-based variability modeling in industry. In *MODELS*, 2014.

7 Thorsten Berger, Ralf Rublack, Divya Nair, Joanne M. Atlee, Martin Becker, Krzysztof Czarnecki, and Andrzej Wąsowski. A Survey of Variability Modeling in Industrial Practice. In *VaMoS*, 2013.

8 Thorsten Berger, Steven She, Krzysztof Czarnecki, and Andrzej Wąsowski. Feature-to-Code mapping in two large product lines. In *SPLC*, 2010.

9 Thorsten Berger, Steven She, Rafael Lotufo, Andrzej Wąsowski, and Krzysztof Czarnecki. Variability modeling in the real: A perspective from the operating systems domain. In *ASE*, 2010.

10 Thorsten Berger, Steven She, Rafael Lotufo, Andrzej Wąsowski, and Krzysztof Czarnecki. A Study of Variability Models and Languages in the Systems Software Domain. *IEEE Transactions of Software Engineering*, 39(12):1611–1640, 2013.

11 John Businge, Openja Moses, Sarah Nadi, Engineer Bainomugisha, and Thorsten Berger. Clone-based variability management in the Android ecosystem. In *ICSME*, 2018.

12 Paul Clements and Linda Northrop. *Software Product Lines: Practices and Patterns*. Addison-Wesley, Boston, MA, 2001.

13 Reidar Conradi and Bernhard Westfechtel. Version models for software configuration management. *ACM Comput. Surv.*, 30(2):232–282, 1998.

14 Krzysztof Czarnecki and Ulrich W. Eisenecker. *Generative Programming: Methods, Tools, and Applications*. Addison-Wesley, Boston, MA, 2000.

15 Danny Dig, Kashif Manzoor, Ralph Johnson, and Tien N Nguyen. Refactoring-aware configuration management for object-oriented programs. In *ICSE*, 2007.

16 Yael Dubinsky, Julia Rubin, Thorsten Berger, Slawomir Duszynski, Martin Becker, and Krzysztof Czarnecki. An exploratory study of cloning in industrial software product lines. In *CSMR*, 2013.

17 Jacky Estublier, David Leblang, André van der Hoek, Reidar Conradi, Geoffrey Clemm, Walter Tichy, and Darcy Wiborg-Weber. Impact of software engineering research on the practice of software configuration management. *ACM Transactions on Software Engineering and Methodology*, 14(4):383–430, 2005.

18 Kyo Kang, Sholom Cohen, James Hess, William Nowak, and Spencer Peterson. Feature-oriented domain analysis (FODA) feasibility study. Technical Report SEI-90-TR-21, CMU, 1990.

19 Timo Kehrer, Udo Kelter, and Gabriele Taentzer. A rule-based approach to the semantic lifting of model differences in the context of model versioning. In *ASE*, 2011.

20 Jacob Krueger, Wanzi Gu, Hui Shen, Mukelabai Mukelabai, Regina Hebig, and Thorsten Berger. Towards a better understanding of software features and their characteristics: A case study of Marlin. In *VaMoS*, 2018.

21 Vincent J. Kruskal. Managing multi-version programs with an editor. *IBM Journal of Research and Development*, 28(1):74–81, 1984.

22 Philip Langer, Manuel Wimmer, Petra Brosch, Markus Herrmannsdörfer, Martina Seidl, Konrad Wieland, and Gerti Kappel. A posteriori operation detection in evolving software models. *Journal of Systems and Software*, 86(2):551–566, 2013.

23 Lukas Linsbauer, Thorsten Berger, and Paul Grünbacher. A classification of variation control systems. In *GPCE*, 2017.

24 Stephen A. MacKay. The state of the art in concurrent, distributed configuration management. In *SCM-4 and SCM-5*, 1995.

25 Axel Mahler. Configuration management. Chapter Variants: Keeping Things Together and Telling Them Apart. Wiley, 1995.

26 Jean Melo, Claus Brabrand, and Andrzej Wąsowski. How does the degree of variability affect bug finding? In *ICSE*, ACM.

27 Mukelabai Mukelabai, Damir Nešić, Salome Maro, Thorsten Berger, and Jan-Philipp Steghöfer. Tackling combinatorial explosion: A study of industrial needs and practices for analyzing highly configurable systems. In *ASE*, 2018.

28 David Parnas. On the design and development of program families. *IEEE Transactions on Software Engineering*, 2(1):1–9, 1976.

29 Leonardo Passos, Jesus Padilla, Thorsten Berger, Sven Apel, Krzysztof Czarnecki, and Marco Tulio Valente. Feature scattering in the large: A longitudinal study of Linux kernel device drivers. In *MODULARITY*, 2015.

30 Christopher Pietsch, Timo Kehrer, Udo Kelter, Dennis Reuling, and Manuel Ohrndorf. SiPL–A Delta-Based Modeling Framework for Software Product Line Engineering. In *ASE*, 2015.

31 Julio Sincero, Horst Schirmeier, Wolfgang Schröder-Preikschat, and Olaf Spinczyk. Is the Linux kernel a software product line. In *Workshop on Open Source Software and Product Lines*, 2007.

32 Daniel Strueber, Mukelabai Mukelabai, Jacob Krueger, Stefan Fischer, Lukas Linsbauer, Jabier Martinez, and Thorsten Berger. Facing the truth: Benchmarking the techniques for the evolution of variant-rich systems. In *SPLC*, 2019.

33 Christer Thörn. Current state and potential of variability management practices in software-intensive SMEs: Results from a regional industrial survey. *Information and Software Technology*, 52(4):411–421, 2010.

34 Eric Walkingshaw and Klaus Ostermann. Projectional editing of variational software. In *GPCE*, 2014.

35 Bernhard Westfechtel, Bjørn P. Munch, and Reidar Conradi. A layered architecture for uniform version management. *IEEE Transactions of Software Engineering*, 27(12):1111–1133, 2001.

36 Thomas Thüm, Leopoldo Teixeira, Klaus Schmid, Eric Walkingshaw, Mukelabai Mukelabai, Mahsa Varshosaz, Goetz Botterweck, Ina Schaefer, and Timo Kehrer. Towards efficient analysis of variation in time and space. In *VaMoS*, 2019.

37 Sofia Ananieva, Timo Kehrer, Heiko Klare, Anne Koziolek, Henrik Lönn, Ramesh Sethu, Andreas Burger, Gabriele Taentzer, and Bernhard Westfechtel. Towards a conceptual model for unifying variability in space and time. In *VaMOS*, 2019.

6

6.30 Visual Analytics for Sets over Time and Space

Organizers: Sara Irina Fabrikant, Silvia Miksch, and Alexander Wolff

Seminar No. 19192

Date: May 5–10, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.5.31

Participants: Hugo Akitaya, Gennady Andrienko, Natalia V. Andrienko, Michael Bekos, Susanne Bleisch, Steven Chaplick, Sabine Cornelsen, Somayeh Dodge, William Evans, Sara Irina Fabrikant, Fabian Frank, Panos Giannopoulos, Jan-Henrik Haunert, Michael Kaufmann, Philipp Kindermann, Stephen G. Kobourov, Giuseppe Liotta, Eva Mayr, Tamara Mchedlidze, Wouter Meulemans, Silvia Miksch, Martin Nöllenburg, Yoshio Okamoto, Margit Pohl, Peter Rodgers, André Schulz, André Skupin, Christian Tominski, Alexander Wolff

Seminar Goals

Increasing amounts of data offer great opportunities to promote technological progress and business success. Visual analytics aims at enabling the exploration and the understanding of large and complex data sets by intertwining interactive visualization, data analysis, human-computer interaction, as well as cognitive and perceptual science. Cartography has for thousands of years dealt with the depiction of spatial data, and more recently geovisual analytics researchers have joined forces with the visual analytics community to create visualizations to help people to make better and faster decisions about complex problems that require the analysis of big data.

Set systems comprise a generic data model for families of sets. A *set* is defined as a collection of unique objects, called the set elements, with attributes, membership functions, and rules. Such a complex data model asks for appropriate exploration methods. As with many types of data, set systems can vary over time and space. It is important, however, not to treat time and space as usual variables. Their special characteristics such as different granularities, time primitives (time points vs. intervals), hierarchies of geographic or administrative regions need to be taken into account. Visualizing and analyzing such changes is challenging due to the size and complexity of the data sets.

Sets systems can also be seen as hypergraphs where the vertices represent the ground elements and the edges are the sets. However, compared to conventional graphs that represent only binary relations (that is, sets with two elements), the visualization of general hypergraphs has received little attention. This is even more so when dealing with dynamic hypergraphs or hypergraphs that represent spatial information.

In this seminar, we aimed at bringing together researchers from the areas of visual analytics, information visualization and graph drawing, geography and GIScience, as well as cartography and (spatial) cognition, in order to develop a theory and visualization methods for set systems that vary over time and space.

Seminar Program

As the topic of the seminar was interdisciplinary and the participants had very different scientific backgrounds, we introduced the main themes of the seminar in three separate sections: "Sets in Time", "Sets in Space", and "Graph Drawing and Set Visualization". Each section consisted of a survey talk and three to four short talks. The three sections were followed by a panel discussion. For the survey talks, we explicitly asked the presenters to give a balanced overview over their area (rather than to focus on their own scientific contributions).

On the second day of the seminar, we collected a number of challenging open problems. Then we formed five groups, each of which worked on a specific open problem for the remainder of the seminar. The work within the groups was interrupted only a few times; in order to share progress reports, listen to open-mic talks, and to discuss possible future activities. These plenary meetings helped to exchange the different visions of the working groups.

We now list the items of the program in detail.

1. Section "Sets in Time" (for abstracts, see Section 3 of the full report)
   - Peter Rogers gave an excellent survey talk about techniques for visualizing sets over time. He illustrated possible challenges and opportunities in this research area.
   - Philipp Kindermann presented some results and open questions in simultaneous orthogonal graph drawing.
   - Wouter Meuleman talked about spatially and temporally coherent visual summaries.
   - Tamara Mchedlidze introduced a data-driven approach to quality metrics of graph visualizations.

- Margit Pohl discussed perception considerations of space and time in cognitive psychology and their implications for the design of visualizations.

2. Section "Sets in Space" (for abstracts, see Section 4 of the full report)
   - Sara Fabrikant gave an inspiring survey talk about space discussed from a cartographer's view.
   - Natalia Andrienko elaborated about evolving sets in space.
   - Somayeh Dodge discussed dynamic visualization of interaction in movement of sets.
   - Jan-Henrik Haunert introduced fast retrieval of abstracted representations for sets of points within user-specified temporal ranges.
3. Section "Graph Drawing and Set Visualization" (for abstracts, see Section 5 of the full report)
   - André Schulz very nicely surveyed the area of drawing graphs and hypergraphs and sketched the main challenges in this area.
   - Michalis Bekos gave a short overview of graph drawing beyond planarity.
   - Sabine Cornelsen talked about general support for hypergraphs.
   - Martin Nöllenburg introduced plane supports for spatial hypergraphs.
4. The panel discussion was entitled "Visual Analytics for Sets over Time and Space: What are the burning scientific questions? An interdisciplinary perspective." André Skupin, Steven Kobourov, and Susanne Bleisch each gave a short statement about the central questions of his or her area; see Section 6 of the full report. Afterwards we had a fruitful and interesting discussion, which led to a productive open problem session.
5. The working groups formed around the following open problems:
   - "Concentric Set Schematization",
   - "From Linear Diagrams to Interval Graphs",
   - "Thread Visualization",
   - "Clustering Colored Points in the Plane", and
   - "Flexible Visualization of Sets over Time and Space".

The reports of the working groups were collected by Michalis Bekos, Steven Chaplick, William Evans, Jan-Henrik Haunert, and Christian Tominski; see Section 7 of the full report.

Future Plans

During our seminar, plans for a follow-up seminar were discussed in a plenary meeting. The seminar-to-be will aim at integrating the approaches for set visualization that have been taken by the different communities (geovisualization, information visualization, and graph drawing, including industry and research). Susanne Bleisch, Steven Chaplick, Jan-Henrik Haunert, and Eva Mayr are currently discussing the precise focus and a title to match that focus.

Among the 29 participants of the seminar, 24 participated in the survey that Dagstuhl does at the end of every seminar. Many answers were in line with the average reactions that Dagstuhl collected over a period of 60 days before our seminar (such as the scientific quality of the seminar, which received a median of 10 out of 11 – "outstanding"). A few questions, however, received different feedback. For example, due to the interdisciplinary nature of the seminar, we had more frequent Dagstuhl visitors than usually: a third of the participants of the survey had been to Dagstuhl at least seven times. It was also interesting to see that more participants than usually stated that our seminar had inspired new research ideas, joint projects or publications, that it had led to insights from neighboring fields, and that it had identified new research directions.

In spite of the organizers' attempt to have a diverse group of participants, all survey participants were from academia and only two rated themselves as "junior". Not surprisingly, some participants suggested to have more PhD students, more people from industry, and generally more people from applications rather than from (graph drawing) theory. The last free text comment in the survey reads: "Once again, a great week at Schloss Dagstuhl – thank you!"

Acknowledgments

We all enjoyed the unique Dagstuhl atmosphere. In particular, it was great to have the opportunity to use a separate room for each working group. We thank Philipp Kindermann for collecting the self-introductory slides before the seminar and for assembling this report after the seminar.

6.31 Approaches and Applications of Inductive Programming

Organizers: Luc De Raedt, Richard Evans, Stephen H. Muggleton, and Ute Schmid

Seminar No. 19202

Date: May 12–17, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.5.58

Participants: Eli Bingham, Ivan Bratko, Maurice Chandoo, Lidia Contreras-Ochando, Andrew Cropper, Luc De Raedt, Amit Dhurandhar, Kevin Ellis, Richard Evans, Cesar Ferri Ramirez, Johannes Fürnkranz, Elena Leah Glassman, José Hernández-Orallo, Céline Hocquette, Tomáš Kliegr, Mark Law, Eneldo Loza Mencía, Stephen H. Muggleton, Stassa Patsantzis, Ruzica Piskac, Johannes Rabold, Harald Ruess, Ute Schmid, Michael Siebers, Armando Solar-Lezama, Gust Verbruggen

Inductive programming 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 [11]. 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 [15, 29, 36, 37]. On the other hand, inductive programming is a special approach to program synthesis. It complements deductive and transformational approaches [4, 25, 39]. 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 [35]. Finally, symbolic approaches can be combined with probabilistic methods [8, 9].

Inductive program synthesis is of interest for researchers in artificial intelligence since the late sixties [2]. 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 [5, 6, 22]. For example, the Microsoft Excel plug-in Flashfill synthesizes programs from a small set of observations of user behavior [13–15]. Related applications are in process mining and in data wrangling [19, 21]. Inductive programming in general offers powerful approaches to learning from relational data [23, 30] and to learning from observations in the context of autonomous intelligent agents [20, 28, 36]. Furthermore, inductive programming can be applied in the context of teaching programming [38, 41].

A recent new domain of interest is how to combine inductive programming with blackbox approaches, especially in the context of (deep) neural networks [10] and in data science.

Relation to Previous Dagstuhl-Seminars

The seminar is a continuation Dagstuhl-Seminars 13502, 15442, and 17382. In the first seminar, the focus was on establishing the research community by exploring the different areas of basic research and applications of inductive programming and identifying commonalities and differences in methods and goals. In the second seminar, more in-depth coverage of algorithmic methods was provided and the relation of inductive programming to cognitive modeling was explored. The third seminar had a main focus on applications in data cleansing, teaching programming, and interactive training. Furthermore, first proposals for neural approaches to learning for inductive programming were presented.

Besides many new insights from many discussions, visible outcomes from the previous seminars are:

- Muggleton, S.H., Schmid, U., Zeller, C., Tamaddoni-Nezhad, A. and T. Besold (2019). Ultra-strong machine learning – comprehensibility of programs learned with ILP. Machine Learning, 107(7), 1119–1140.
- Schmid, U., Zeller, C., Besold, T., Tamaddoni-Nezhad, A., & Muggleton, S.H. (2017). How does predicate invention affect human comprehensibility?. In Alessandra Russo and James Cussens, editors, Proceedings of the 26th International

Conference on Inductive Logic Programming (ILP 2016), pp. 52-67, Springer.
- Hernández-Orallo, J., Martínez-Plumed, F., Schmid, U., Siebers, M., & Dowe, D. L. (2016). Computer models solving intelligence test problems: Progress and implications. Artificial Intelligence, 230, 74-107.
- Gulwani, S., Hernández-Orallo, J., Kitzelmann, E., Muggleton, S. H., Schmid, U., & Zorn, B. (2015). Inductive programming meets the real world. *Communications of the ACM*, 58(11), 90-99.
- https://en.wikipedia.org/wiki/Inductive_programming
- NIPS'2016 workshop on Neural Nets and Program Induction involving Stephen Muggleton.
- A collaboration in the context of the EPSRC funded Human-Like Computing Network+ headed by Muggleton (see http://hlc.doc.ic.ac.uk/).

For the fourth seminar, we extend our invitation to researchers from deep learning and to researchers addressing statistical machine learning and probabilistic reasoning. **The focus of the fourth seminar has been on the potential of inductive programming for explainable AI, especially in combination with (deep) neural networks and with data science.**

Inductive Programming as a Approach to Explainable AI

Recently it has been recognized that explainability is crucial for machine learning to be usable in real world domains – especially such where erroneous decisions might be harmful to humans. Consequently, interfaces which explain (aspects) of classifier decisions have been proposed – especially in the context of deep learning [3, 32]. The notion of explainability has already been addressed in the early days of machine learning: During the 1980s Michie defined Machine Learning in terms of two orthogonal axes of performance: predictive accuracy and comprehensibility of generated hypotheses. Since predictive accuracy was readily measurable and comprehensibility not so, later definitions in the 1990s, such as that of Mitchell [27], tended to use a one-dimensional approach to Machine Learning based solely on predictive accuracy, ultimately favouring statistical over symbolic Machine Learning approaches.

In [29] a definition was provided of comprehensibility of hypotheses which can be estimated using human participant trials. Experiments were conducted testing human comprehensibility of logic programs. Results show that participants were not able to learn the relational concept on their own from a set of examples but they were able to apply the relational definition provided by the ILP Metagol system correctly. That is, the results demonstrate that ILP systems can fulfill Michie's criterion of operational effectiveness. The findings also imply the existence of a class of relational concepts which are hard to acquire for humans, though easy to understand given an abstract explanation. We believe improved understanding of this class could have potential relevance to contexts involving human learning, teaching and verbal interaction.

Finally, while research in explanations in the context of neural networks is focusing on visualization, ILP learned classifiers allow natural language explanations. In the Dagstuhl seminar we plan to discuss possibilities for combining deep learning approaches and inductive programming such that both modes of explanations can be generated.

Inductive Programming for Support in Data Science

The success of the inductive programming system FlashFill has motivated the developed of several approaches to using inductive programming in the context of data science, more specifically data wrangling. It is well known that in data science and data mining processes, about 80 per cent of the time goes to selecting the right data, and further pre-processing it so that it be input into data mining software. One important step in that process is data wrangling, which is concerned with cleaning up the data and transforming it across different formats. For this step, inductive programs can be used; various approaches are moving in that direction, e.g. FlashRelate [1], Tacle [19] and SYNTH [7]. Furthermore, workshops are being organised on the topic of automating data wrangling (e.g. at ICDM 2016, http://dmip.webs.upv.es/DWA2016/, at ECMLPKDD 2019, https://sites.google.com/view/autods and at Dagstuhl Seminar 18401), and tools such as MagicHaskeller and JailBreakR being used in this context. In the Dagstuhl seminar, we also want to deepen the link between data wrangling and inductive programming.

Inductive Programming and Neural Computation

The deep learning community has been interested in taking on challenging tasks from the artificial intelligence community. It is therefore no surprise that they have also started to look into inductive and automatic programming. In particular, they have contributed several mixtures of traditional computational models with those of neural networks. For instance, the neural Turing machine [12] integrates a neural network with an external memory and it is able to learn simple algorithms such as copy and sort, the neural program interpreter [31] is a recurrent neural network that learns to represent and execute programs from program traces, while [33] present an end-to-end differentiable interpreter for the programming language Forth and [24, 34] for a declarative Prolog like language. The central goal in these approaches is to obtain an end-to-end differentiable model. While initial results are promising, the approaches still require a lot of data to be trained or need to scale up. This contrasts with the more traditional symbolic approaches to inductive programming and thus a key opportunity is to further cross-fertalize these two approaches.

Inductive Programming and Human-like Computing

The human ability to master complex demands is to a large extend based on the ability to exploit previous experiences. Based on our experience, we are able to predict characteristics or reactions of (natural or man-made, inanimate or animate) objects, we can reason about possible outcomes of actions, and we can apply previously successful routines and strategies to new tasks and problems. In philosophy, psychology and artificial intelligence, researchers proposed that the core process to expand knowledge, that is, construct hypotheses, in such a way that we can transfer knowledge from previous experience to new situations is *inductive* inference [16, 18, 40].

One special aspect of induction is that humans are able to acquire complex, productive rule sets from experience. Following Chomsky, rules are productive when they can be applied in situations of various complexity. Typical examples of such rule sets are knowledge about natural language grammar, recursive concepts such as ancestor and recursive problem solving strategies. For

example, if humans have learned how to solve Tower of Hanoi problems with three and four discs, at least some of them are able to generalize the underlying strategy for solving problems with an arbitrary number of discs.

Inductive programming provides mechanisms to generate such productive, recursive rule sets. Examples of recent work on using inductive programming to model this learning-capability of human cognition are [17, 20, 23, 26, 36]. Therefore, it might be fruitful for cognitive scientists to get acquainted with inductive programming as one approach to model the acquisition of complex knowledge structures. On the other hand, knowledge gained from experiments in human problem solving, concept learning and language acquisition can be a source of inspiration for new algorithmic approaches to inductive programming.

Objectives and Expected Outcomes of the Seminar

A long-term objective of the seminar series is to establish inductive programming as a self-contained research topic in artificial intelligence, especially as a field of machine learning and of cognitive modeling. The seminar serves as community building event by bringing together researchers from different areas of inductive programming – especially inductive logic programming and inductive functional programming –, from different application areas such as end-user programming and tutoring, and from cognitive science research, especially from cognitive models of inductive (concept) learning. For successful community building we seek to balance junior and senior researchers and to mix researchers from universities and from industry.

The previous seminars resulted in new collaborations between researchers from different backgrounds as documented in joint publications and we expect that the collaborations will continue, deepen and extend, resulting not only in further joint publications but also in joint research projects.

In the fourth seminar, we continued and extended previous discussions addressing 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 are more suited than standard machine learning approaches, including deep learning and probabilistic programming. Focus here is on possibilities of combining (deep) neural approaches or probabilistic programming with (symbolic) inductive programming, especially with respect to new approaches to comprehensibility of machine learned models and on explainable AI.
- 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 end-user programming and programming education and identifying further relevant areas of application.
- Establishing stronger relations between cognitive science research on inductive learning and inductive programming under the label of human-like computation.
- Strengthening the relation of inductive programming and data science, especially with respect to data cleansing and data wrangling.

Concluding Remarks and Future Plans

In the wrapping-up section, we decided to move the IP webpage[51] to a Wiki and encouraged all participants to make available their systems, tutorial/lecture slides and publications there.

As the grand IP challenge we came up with 2017 is still up:

> 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 Daniel W. Barowy, Sumit Gulwani, Ted Hart, and Benjamin Zorn. Flashrelate: Extracting relational data from semi-structured spreadsheets using examples. *SIGPLAN Not.*, 50(6):218–228, June 2015.

2 A. W. Biermann, G. Guiho, and Y. Kodratoff, editors. *Automatic Program Construction Techniques*. Macmillan, New York, 1984.

3 Alexander Binder, Sebastian Bach, Gregoire Montavon, Klaus-Robert Müller, and Wojciech Samek. Layer-wise relevance propagation for deep neural network architectures. In *Information Science and Applications (ICISA) 2016*, pages 913–922. Springer, 2016.

4 Rastislav Bodik and Emina Torlak. Synthesizing programs with constraint solvers. In *CAV*, page 3, 2012.

5 A. Cypher, editor. *Watch What I Do: Programming by Demonstration*. MIT Press, Cambridge, MA, 1993.

6 Allen Cypher, Mira Dontcheva, Tessa Lau, and Jeffrey Nichols, editors. *No Code Required: Giving Users Tools to Transform the Web*. Elsevier, 2010.

7 Luc De Raedt, Hendrik Blockeel, Samuel Kolb, Stefano Teso, and Gust Verbruggen. Elements of an automatic data scientist. In *International Symposium on Intelligent Data Analysis*, pages 3–14. Springer, 2018.

8 Luc De Raedt and Angelika Kimmig. Probabilistic (logic) programming concepts. *Machine Learning*, 100(1):5–47, 2015.

9 Richard Evans and Edward Grefenstette. Learning explanatory rules from noisy data. *Journal of Artificial Intelligence Research*, 61:1–64, 2018.

10 Richard Evans, David Saxton, David Amos, Pushmeet Kohli, and Edward Grefenstette. Can neural networks understand logical entailment? *arXiv preprint arXiv:1802.08535*, 2018.

11 P. Flener and U. Schmid. Inductive programming. In C. Sammut and G. Webb, editors, *Encyclopedia of Machine Learning*, pages 537–544. Springer, 2010.

12 Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *CoRR*, abs/1410.5401, 2014.

13 Sumit Gulwani. Automating string processing in spreadsheets using input-output examples. In *38th Symposium on Principles of Programming Languages*. ACM, 2011.

[51] www.inductive-programming.org

14 Sumit Gulwani, William R. Harris, and Rishabh Singh. Spreadsheet data manipulation using examples. *Communications of the ACM*, 55(8):97–105, 2012.

15 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.

16 Ulrike Hahn, Todd M Bailey, and Lucy BC Elvin. Effects of category diversity on learning, memory, and generalization. *Memory & Cognition*, 33(2):289–302, 2005.

17 J. Hernández-Orallo, D. L. Dowe, and M. V. Hernández-Lloreda. Universal psychometrics: Measuring cognitive abilities in the machine kingdom. *Cognitive Systems Research*, 27(0):50–74, 2014.

18 J.H. Holland, K.J. Holyoak, R.E. Nisbett, and P.R. Thagard. *Induction – Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA, 1986.

19 Samuel Kolb, Sergey Paramonov, Tias Guns, and Luc De Raedt. Learning constraints in spreadsheets and tabular data. *Machine Learning*, 106(9-10):1441–1468, 2017.

20 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.

21 Vu Le and Sumit Gulwani. Flashextract: A framework for data extraction by examples. *ACM SIGPLAN Notices*, 49(6):542–553, 2014.

22 Henry Lieberman, editor. *Your Wish is My Command: Programming by Example*. Morgan Kaufmann, San Francisco, 2001.

23 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.

24 Robin Manhaeve, Sebastijan Dumancic, Angelika Kimmig, Thomas Demeester, and Luc De Raedt. Deepproblog: Neural probabilistic logic programming. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 3749–3759. Curran Associates, Inc., 2018.

25 Zohar Manna and Richard Waldinger. A deductive approach to program synthesis. *ACM Transactions on Programming Languages and Systems*, 2(1):90–121, 1980.

26 G. F. Marcus. *The Algebraic Mind. Integrating Conncetionism and Cognitive Science*. Bradford, Cambridge, MA, 2001.

27 Tom Mitchell. *Machine learning*. McGraw Hill, 1997.

28 Bogdan Moldovan, Plinio Moreno, Davide Nitti, José Santos-Victor, and Luc De Raedt. Relational affordances for multiple-object manipulation. *Auton. Robots*, 42(1):19–44, 2018.

29 S.H. Muggleton, U. Schmid, C. Zeller, A. Tamaddoni-Nezhad, and T. Besold. Ultra-strong machine learning – comprehensibility of programs learned with ILP. *Machine Learning*, 2018.

30 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.

31 Scott E. Reed and Nando de Freitas. Neural programmer-interpreters. *CoRR*, abs/1511.06279, 2015.

32 Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1135–1144. ACM, 2016.

33 Sebastian Riedel, Matko Bosnjak, and Tim Rocktäschel. Programming with a differentiable forth interpreter. *CoRR*, abs/1605.06640, 2016.

34 Tim Rocktäschel and Sebastian Riedel. End-to-end differentiable proving. *CoRR*, abs/1705.11040, 2017.

35 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.

36 Ute Schmid and Emanuel Kitzelmann. Inductive rule learning on the knowledge level. *Cognitive Systems Research*, 12(3):237–248, 2011.

37 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.

38 Rishabh Singh, Sumit Gulwani, and Armando Solar-Lezama. Automated feedback generation for introductory programming assignments. *ACM SIGPLAN Notices*, 48(6):15–26, 2013.

39 Douglas R. Smith. The synthesis of LISP programs from examples: A survey. In *Automatic Program Construction Techniques*, pages 307–324. Macmillan, 1984.

40 J. Tenenbaum, T.L. Griffiths, and C. Kemp. Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, 10(7):309–318, 2006.

41 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

6.32 Enumeration in Data Management

Organizers: Endre Boros, Benny Kimelfeld, Reinhard Pichler, and Nicole Schweikardt

Seminar No. 19211

Date: May 19–24, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.5.89

Participants: Amir Abboud, Kira V. Adaricheva, Antoine Amarilli, Kristof Berczi, Christoph Berkholz, Endre Boros, Pierre Bourhis, Nofar Carmeli, Ondrej Cepek, Nadia Creignou, Arnaud Durand, Khaled M. Elbassioni, Etienne Grandjean, Alejandro J. Grez, Aritanan Gruber, Mamadou Moustapha Kanté, Batya Kenig, Benny Kimelfeld, Christoph Koch, Phokion G. Kolaitis, Markus Kröll, Ester Livshits, Kazuhisa Makino, Andrea Marino, Wim Martens, Stefan Mengel, Shin-Ichi Nakano, Matthias Niewerth, Lhouari Nourine, Liat Peterfreund, Reinhard Pichler, Cristian Riveros, Yehoshua Sagiv, Nicole Schweikardt, Thomas Schwentick, Luc Segoufin, Yann Strozecki, Alexandre Termier, Etsuji Tomita, György Turan, Martin Ugarte, Takeaki Uno, Stijn Vansummeren, Alexandre Vigny, Heribert Vollmer, Thomas Zeume

In recent years, various concepts of enumeration have arisen in the fields of Databases, Computational Logic, and Algorithms, motivated by applications of data analysis and query evaluation. Common to all concepts is the desire to compute a stream of items with as small as possible waiting time between consecutive items, referred to as the "delay." Alongside each concept, there evolved algorithmic techniques for developing solvers, and proof techniques for establishing complexity bounds. In addition to the traditional guarantees of "polynomial delay" and "incremental polynomial," researchers have been pursuing stronger guarantees such as "constant delay" in the context of logical query evaluation, "dynamic complexity" of incremental maintenance, and "factorized databases." The growing interest and rapid evolution of the associated research brings up opportunities of significantly accelerating the computation of big results, by devising and adopting general-purpose methodologies.

In Dagstuhl Seminar 19211 on "Enumeration in Data Management," key researchers from relevant communities have gathered to gain a better understanding the recent developments, lay out the important open problems, and join forces towards solutions thereof. These communities include researchers who explore enumeration problems in the fields of *databases*, *logic*, *algorithms* and *computational complexity*. We have had invited tutorials by

- Luc Segoufin on *Constant-delay enumeration*
- Takeaki Uno on *Enumeration algorithms*
- Yann Strozecki on *Enumeration complexity – defining tractability*
- Markus Kröll on *Enumeration complexity – a complexity theory for hard enumeration problems*
- Endre Boros on *Monotone generation problems*.

We also had presentations by most of the other participants. Moreover, the participants have prepared in advance a list of open problems in a document that we shared and jointly maintained. We have discussed the open problems during designated times of the seminar.

The organizers are highly satisfied with the seminar. We have got a very high acceptance rate for our invitations. In fact, there were further researchers whom we would have liked to invite after the first invitation round but, unfortunately, no room was left. The participants were exceptionally involved and engaged. Some considerable progress has been made on the open problems prepared in advance, as will be reported in future publications that will acknowledge the seminar. The seminar has also initiated joint efforts to disseminate toolkits for data-centric enumeration problems, including algorithmic techniques, proof techniques, and important indicator problems. To this end, we have had sessions of working groups for the different types of toolkit components. In particular, we have initiated a Wikipedia page on enumeration algorithms:

https://en.wikipedia.org/wiki/Enumeration_algorithm

This page will evolve to contain a thorough picture of the principles and techniques of enumeration problems.

6.33 Topology, Computation and Data Analysis

Organizers: Michael Kerber, Vijay Natarajan, and Bei Wang

Seminar No. 19212

Date: May 19–24, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.5.110

Participants: Ulrich Bauer, Ranita Biswas, Georges-Pierre Bonneau, Roxana Bujack, Hamish Carr, Tamal K. Dey, Pawel Dlotko, Oliver Gäfvert, Christoph Garth, Hans Hagen, Ingrid Hotz, Michael Kerber, Claudia Landi, Michael Lesnick, Vijay Natarajan, Arnur Nigmetov, Emilie Purvine, Vanessa Robins, Paul Rosen, Filip Sadlo, Primoz Skraba, Julien Tierny, Hubert Wagner, Bei Wang, Gunther H. Weber, Eugene Zhang, Yue Zhang

The Dagstuhl Seminar titled "Topology, Computation, and Data Analysis" brought together researchers in mathematics, computer science, and visualization to engage in active discussions on theoretical, computational, practical, and application aspects of topology for data analysis. The seminar has led to stronger ties between the computational topology and TopoInVis (topology based visualization) communities and identification of research challenges and open problems that can be addressed together.

Context

Topology is the study of connectivity of space that abstracts away geometry and provides succinct representations of the space and functions defined on it. Topology-based methods for data analysis have received considerable attention in the recent years given its promise to handle large and feature-rich data that are becoming increasingly common. Computing topological properties in the data domain and/or range is a step in the direction of more abstract, higher-level data analysis and visualization. Such an approach has become more important in the context of automatic and semi-automatic data exploration, analysis, and understanding. The primary attraction for topology-based methods is the ability to generate "summary" qualitative views of large data sets. Such views often require fewer geometrical primitives to be extracted, stored, and to be visualized as compared to views obtained directly from the raw data. Two communities, computational topology and TopoInVis (topology based visualization), have made significant progress during the past two decades on developing topological abstractions and applying them to data analysis. In addition, there are multiple other research programs (relatively fewer in number) on this topic within the statistics and machine learning fields, and within a few application domains. Computational topology grew from within computational geometry and algebraic topology and studies algorithmic questions on topological structures. The focus of topological data analysis and TopoInVis is data – algorithms, methods, and systems for improved and intuitive understanding of data via application of topological structures. Researchers in computational topology typically have a math or theoretical computer science background whereas TopoInVis researchers have a computational, computer engineering, or applied background. There is very little communication between the two communities due to the different origins and the fact that there are no common conferences or symposia where both communities participate.

Goals

The Dagstuhl seminar 17292 (July 2017) successfully brought together researchers with mixed background to talk about problems of mutual interest. Following this seminar, the benefits of the inter-community ties was well appreciated, at least by the attendees of the seminar. The goal of the current seminar was to strengthen existing ties, establish new ones, identify challenges that requires the two communities to work together, and establish mechanisms for increased communication and transfer of results from one to the other. During the previous Dagstuhl seminar, we also noticed significant interaction between researchers within the individual communities, with say theoretical and applied backgrounds. We wanted to continue to encourage such interaction.

Topics

We chose four current and emerging topics that will benefit from an inter-community discussion. Topics are common to both communities, with different aspects studied within an individual community.

Reeb graphs, Reeb Spaces, and Mappers. The Reeb graph, its loop-free version called the contour tree, and the higher-dimensional generalization called the Reeb space are topological structures that capture the connectivity of level sets of univariate or multivariate functions. They are independently well studied within the computational topology and TopoInVis communities. Recent developments define stable distance measures between Reeb graphs, inspired by analogous distance measures in persistent homology. Barring a few exceptions, the theoretical results have no practical realizations. On the practical side, effective visual exploration and visual analysis methods based on Reeb graphs and spaces have been developed for a wide variety of domains including combustion studies, climate science, astronomy, and molecular modeling. These applications often utilize only a simplified version of the topological structure. One such simplification, the mapper algorithm, consists of a discretized version of Reeb graphs and has shown an immense industrial potential. Very recently, the theoretical aspects of the mapper algorithm and its generalizations has moved in the focus of research. Exchange of ideas and results between the two communities will help advancing this progress further.

Topological analysis and visualization of multivariate data. Multivariate datasets arise in many scientific applications. Consider, for example, combustion or climate simulations where multiple physical measurements (say, temperature and pressure) or concentrations of chemical species are computed simultaneously. We model these variables mathematically as multiple continuous, real-valued functions. We are interested in understanding the relationships between these functions, and more generally, in developing efficient and effective tools for their analysis and visualization. Unlike for real-valued functions, very few tools exist for studying multivariate data topologically. Besides the aforementioned Reeb spaces and mappers, notable examples of these tools are the Jacobi sets, Pareto sets, and Joint Contour Nets. Understanding the theoretical properties of these tools and adapting them in analysis and visualization remains a very active research area. In addition, combining these topological tools with multivariate statistical analysis would be of interest. On the other hand, research towards multidimensional persistence would help advance multivariate data analysis both mathematically and computationally. We plan to expand our discussion on multidimensional persistent homology that include topics such as identifying meaningful and computable topological invariants; discussing computability and applicability in the multidimensional setting, comparison of multidimensional data, kernel methods for multidimensional persistence, and adapting multidimensional persistence in visualization.

New opportunities for vector field topology. Vector field topology for visualization pioneered by Helman and Hesselink has inspired much research in topological analysis and visualization of vector fields. A large body of work for time-independent vector field deals with fixed (critical) points, invariant sets, separatrices, periodic orbits, saddle connectors and Morse decomposition as well as vector field simplification that reduces its complexity. Research for time-dependent vector field is concerned with critical point tracking, Finite Time Lyapunov Exponents (FTLE), Lagrangian coherent structure (LCS), streak line topology, as well as unsteady vector field topology. For this workshop, we ask the following questions: can advancements in computational topology help bring new opportunities for the study of vector field topology? In particular, can they help developing novel, scalable and mathematically rigorous ways to rethink vector field data? An example is the topological notion of robustness, a cousin of persistence, introduced via the well diagram and well group theory. Robustness has been shown to be very useful in quantifying feature stability for steady and unsteady vector fields.

Software tools and libraries. How do we make topological data analysis applicable to large datasets? A natural first step is algorithm and software engineering. This refers to developing the best algorithms for a particular problem and to optimize the implementation of these algorithms. The state of affairs within the communities is quite diverse: while scalable algorithms are available for some problems(e.g., computation of Reeb graphs or persistence diagrams in low dimensions), current developments make significant progress on other fronts, for example the computation of approximate persistence diagrams of Vietoris-Rips complexes. On the other extreme, the theory of multi-dimensional persistence is just beginning to be supported by algorithmic contributions. Besides these efforts, parallelizable and distributed algorithms play an important role towards practicality. One further important aspect of software design is interface design, that is, to make those implementations available to non-experts. While this final development step is usually rather neglected in theoretical research, there have been efforts in both communities towards generally applicable and easy-to-use software. Software contributors of both communities will profit from exchanging ideas and experiences.

Participants, Schedule, and Organization

The invitees were identified according to the focus topics of the seminar while ensuring diversity in terms of gender, country / region of workplace, and experience. The aim was to bring together sufficient number of experts interested in each topic and representing the two communities to facilitate an engaging discussion.

We planned for different talk types, longer overviews and shorted contributed research talks, and breakout sessions. We scheduled six overview talks on the first day. These overview talks were aligned with the four topics of the seminar, planned to be accessible to members of both communities, and set the stage for the discussions and shorter research talks on the following days. The speakers Ulrich Bauer (Reeb graphs), Christoph Garth (topology based methods in visualization), Gunther Weber (topological analysis for exascale), Michael Lesnick (computational aspects of 2-parameter persistence) Claudia Landi (multi-parameter persistence), and Vanessa Robins (discrete Morse theory and image analysis) gave a gentle introduction to the area followed by a state-of-the-art report and discussion on open problems.

Participants gave short research talks (16 total) during Tuesday-Friday with a focus on challenges and opportunities. These talks were organized during the morning sessions.

We scheduled breakout sessions on the afternoons of Tuesday and Thursday. On Tuesday, we solicited discussion topics and identified three topics to be of interest – *multivariate data, reconstruction,* and *tensor field topology*. Participants chose to join a group based on their interest. All groups contained participants from both communities. We formed two discussion groups on Thursday. The first group wanted to further discuss *multivariate data* with inputs from experts on multi-parameter persistence who were part of a different group on Tuesday. The second breakout session was on *Multi-parameter persistence computation*, where

they discussed and analyzed a recently proposed algorithm. All groups presented a summary of their discussion and plans during a plenary session at the end of the day.

Many participants joined an organized excursion to Bernkastel-Kues on Wednesday afternoon. On Friday morning, we scheduled a discussion and brainstorming session to close the seminar and and to plan for future events.

Results and Reflection

Participants unanimously agreed that the seminar was successful in enabling cross-fertilization and identifying important challenging problems that require both communities to work together. The breakout sessions were instrumental in identifying some of the challenges and topics for further collaboration. At least two such challenges (together with motivating applications) were identified, possibly leading to collaboratory efforts.

The breakout sessions were planned for the entire afternoon after lunch. The longer duration allowed for in-depth and technical discussions that stimulates further work after the seminar. Based on feedback during informal discussions and the brainstorming session on Friday, we expect multiple working groups will be formed to write expository articles and survey articles. Members of the two communities have also shown enthusiasm to participate in workshops and conferences of each other. In conclusion, we believe that the seminar has achieved the goal of bringing together the two communities and charting a path for tackling bigger challenges in the area of topological data analysis.

6

6.34 Control of Networked Cyber-Physical Systems

Organizers: John S. Baras, Sandra Hirche, Kay Römer, and Klaus Wehrle

Seminar No. 19222

Date: May 26–29, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.5.132

Participants: John S. Baras, Sankar Basu, Marcel Carsten Baunach, Carlos Canudas-de-Wit, Georg Carle, Aaron Ding, Rolf Findeisen, Hannes Frey, René Glebke, James Gross, Andrei Gurtov, Tobias Heer, Thorsten Herfet, Sandra Hirche, Wolfgang Kellerer, Na Li, Mingyan Liu, Mohammad Hossein Mamduhi, Adam Molin, Ehsan Nekouei, Chrysa Papagianni, Daniel Quevedo, Kay Römer, Wolfgang Schröder-Preikschat, Olaf Stursberg, Sebastian Trimpe, Klaus Wehrle, Herbert Werner, Gerhard Wunder, Marco Zimmerling, Martina Zitterbart

Motivation and Purpose of the Seminar

Manufacturing cells and factories, transportation systems and various other parts of critical infrastructure such as energy grids have traditionally been controlled via self-contained, centralized systems continuously monitored and reconfigured by humans. The ever-growing complexity and integration of these Cyber-Physical Systems (CPS) into reconfigurable value chains ("Industrie 4.0"), autonomous cars and other services with high reliability requirements necessitates a radical change in the control strategy: Classic controllers will not be able to handle the massive amounts of data generated by these emerging systems, not only because of restrictions with regard to computational power and complexities that might bar human interventions in the processes, but also due to missing or inadequate methods for the control and the interconnection of the devices comprising such systems. Whilst CPS have moderate bandwidth/throughput requirements, often in the range of a few bytes per control or sensor message, they require high delivery success rates and predictable latency bounds for these messages and the computations performed on the data, often in the order of a few milliseconds. Stable controllers can only be developed if a predictable behavior of the communication and computation infrastructure may be assumed. Otherwise, the systems may not reach the desired states or even become unstable, up to the point where they may cause physical injuries or the loss of human life. Hence, a paradigm shift towards real-time oriented communication and computation in CPSs is necessary.

Such a shift can, however, only be achieved by overcoming the traditionally loose coupling in the design of system components in networks. Currently, both the communication systems community and the control systems community consider the components of the respective other field as a "black box" and abstract from the variations. Valuable insights that the other domain might provide towards the joint goal of keeping a CPS controllable may hence not be available. Although solutions have already been developed that bring communication and control closer together for specific use-cases, the abstraction problem has not been approached from a general, overarching perspective.

The purpose of this seminar was hence to bring together experts working in the key communities relevant for the science of CPSs and Cyber-Physical Networking (CPN) to get a clearer and more detailed picture of the most important issues of the control and networking aspects that CPSs/CPNs bear and to identify the mutual relations and influences of the associated fields, in order to overcome the so-far strict abstractions and boundaries that exist, and to sketch a roadmap for further research in the field. The driving question was how it is possible to derive generalizable co-design methods and metrics that support the development of universal networked CPSs/CPNs.

Prior Dagstuhl Seminars have already addressed CPS aspects such as synthesis (Seminar 17201) and verification methods (Seminar 14122), robustness (Seminar 16362), as well as software engineering for control (Seminar 14382), yet none of these have focused on the interaction, interdependencies and the co-design of communication and control.

Participants and Structure

The seminar brought together a total of 30 participants from various fields within the communication and control domains, ranging from promising young scientists to leading authorities within their respective fields, but also including practitioners from industry with a strong research background, as well as representatives from funding organizations.

The first day of the seminar was dedicated to an in-depth introductory session. Besides as short personal introduction with background and current research interests, **each participant was asked to prepare a personal statement answering the following questions:**

- What are the most important problems to solve in the realm of CPS/CPN?
- What are the main scientific challenges and which fields can contribute to them?
- What have we achieved so far, and what are the pitfalls of past and current research?

Each personal statement was followed by a discussion round on the presented individual statements. The statements and discussions proved highly fruitful, as they allowed the organizers and the participants to gain an understanding of the current state and future challenges in the Control of Networked CPS from the different disciplinary perspectives.

Most often, opinions revolved around the need to understand more about the implications of the dynamic behavior of both the controlled systems themselves and of the communication networks. Research so far seems to have primarily focused on the "steady state", as participant termed it. The uncertainties introduced by controlled systems and (especially wireless) networks in coexistence with other systems, however, seem to call for various improvements in CPS/CPN design. Yet, as other participants expressed it, besides having fostered a better understanding of the basics of the respective other fields in recent years by programs such as DFG's Priority Programme 1914 *Cyber-Physical Networking*, "little" has been achieved be community so far, with a major pitfall being "lopsided" methods which are often attributed to "sticking to domain-specific models". Opening these models to incorporate knowledge from other domains, therefore, seems to be a major challenge for the upcoming time.

A further major topic discussed was the need for more realistic and relevant problem settings in the research efforts, since, as one participant put it, "real problems are more complex than a single inverted pendulum". Hence, to avoid "esoteric" research and thus "ending up as an academic field with zero practical impact", CPS/CPN is in the need of "prov[ing] that what we develop is useful/needed" within the upcoming years. This does not mean that basic research has or needs to be concluded in any way. Yet, further opinions voiced more than once regarded energy efficiency and usable abstraction/decomposition methods (which may at times even sacrifice optimality for applicability and efficiency) as interesting research challenges for the upcoming years, which shows that the community has already begun tackling more practical issues recently. A variety of additional comments showed that few, if any, of the issues of CPS/CPN can be considered as solved by today.

Plenary Discussion: Properties of Cyber-Physical Networks

The unexpected intensity of the discussions following the respective personal introductions revealed the extreme variety of opinions on the nature of CPS/CPN and the major challenges in this interdisciplinary field. To facilitate a common understanding, the personal introductions were thus followed by a plenary discussion on which properties define CPS/CPN and make them interesting for scientific study.

It was agreed that – besides the eponymous intertwining of control, networking and the physically tangible world – CPS/CPN are dominated by *uncertainties* of both the systems and their operational environments, *dynamics* of configuration and load, (usually) *limitations* e.g., with respect to the capacity of the network, computation power and energy, a *control objective* that is sought to achieve through the network (if it is not serving pure monitoring purposes), as well as the associated relative administrative and technical *autonomy* of CPN compared to their traditional counterparts. Regarding typical metrics of timing and scale, it was further agreed that **traditional complexity metrics do not apply to CPN**. There often exist intricate and counterintuitive relationships between timing constraints of control and the network, leading to situations in which certain upper- and lower(!)-bounded delays may even be beneficial for the simplification and stabilization of control. Hence, defining the time-criticality of a system is scenario-dependent. Likewise, scaling effects may lead to situations in which too many local observations may prove counterproductive to controllability so that, depending on the scenario at hand, issues arise regarding the "right" amount of information sharing between local and global players in distributed decision-making processes. As such, **conceiving widely-applicable categories for the complexity of CPS/CPN was identified as an open problem**.

Impulse Talks & Plenary Discussions

For the remaining one and a half days of the seminar, the participants were asked to propose impulse talks on topics related to their respective areas of control of CPS/CPN research. Each talk served as the basis for a subsequent plenary discussion aimed at identifying worthwhile research directions for the community. **Out of a total of 18 proposed impulse talks, six talks were selected by the organizers.** In the following, we present the major insights from the talks and the discussions.

- The development of next-generation wireless communication technologies such as 5G and the increased efficiency of small-scale mobile devices in general, have fueled the interconnection of ever more devices into large-scale CPS. However, as the number of devices generating data and potentially taking action increase, so do the burdens on controllers and the network. In his talk, Carlos Canudas-de-Wit showed first results pointing at the fact that both state estimation and control may provide sufficient results even when only considering a well-chosen aggregating subset of a system's sensing and actuating nodes, as long as the distribution of these nodes follows a specific structure, which can, however, be found for many real-world scenarios. Together with another technique based on partial differential equations, the results of his work showed that when combining both control- and information-/network theoretic models, as well as upcoming techniques such as in-network computation that may provide the necessary aggregation infrastructure, even systems of immense complexity can be controlled without overloading controllers and networks.
- Another challenge of CPS arises when safety guarantees need to be fulfilled, especially when a failure to meet these guarantees can lead to injury or endangerment of human life. Adam Molin presented an industry perspective on the validation and verification of (increasingly) autonomous vehicles, a field in which scenario-based testing approaches represent the state-of-the-art. While the determinism of systems without humans in the loop may aid in the construction of such scenarios, only probabilistic guarantees can be given when humans are involved in the operation of a system. This fact reflects not just on automotives but on multiple other scenarios discussed in the seminar and highlights the importance of joint analysis methods for the control and the communication components of such systems.
- In her talk on 5G Service Automation, Chrysa Papagianni expressed the view that upcoming mobile networks will witness a shift from open-loop to hierarchical closed-loop control as customers shift towards a pay-per-use scheme

for the offered services. Whether the control problems (e.g., regarding network slicing) can be considered to exhibit sub-minute or even real-time requirements (as witnessed in most other systems discussed in the seminar) is still an open question. Yet, considering the anticipated, wide-spread application scenarios of 5G also in the area of CPS/CPN, the seminar identified the issue of base station multi-tenancy as an area for future research within the context of CPN.

- A cornerstone for the successful operation of CPS/CPN are easily-calculable metrics to assess the operational status, as well as to guide the generation, transmission and evaluation of signals within the systems. Vahid Mamduhi in his talk showed that simple age-of-information (AoI) – a common metric applied both by the control and the communications communities in theoretical and practical scenarios – bases on assumptions that can hardly be met by the systems. As a consequence, AoI needs to be augmented by notions of state, timing constraints of the system, and the objective of the control function (all related to a single piece of information) to really provide benefits. Such metrics are arguably hard to conceive for the general case, yet the talk inspired discussions among the participants regarding sensible metrics with broader applicability.
- From a more communication-oriented perspective, James Gross presented his group's efforts towards determining latency bounds in wireless CPS/CPN. Both a queuing-theoretic and a model checking-based approach (the latter concerning a practical implementation of an ultra-reliable low latency protocol) yielded qualitative results that seem promising. Yet, the practical applicability of such approaches is currently hampered by assumptions regarding distortion that may not hold in practice. In the subsequent discussion, topics included *(a)* the question whether making the network completely deterministic (or the ability to make determinism assumptions) is actually needed and achievable, and how possible compromises may look like, *(b)* to which degree techniques such as software-defined networking, in-network processing, time slicing and standards such as 5G can contribute towards such goals, and *(c)* which interfaces, abstractions and design patterns should exist that allow specifying and proving certain guarantees in CPN, especially regarding the interplay of control algorithms and networks.
- The complexity and variety of communication protocols within automation is addressed by the recent Time-Sensitive Networking (TSN) efforts of the IEEE, which seek to offer a vendor-neutral Ethernet-based solution catering both legacy and future real-time applications, including control. Eventually, the automation pyramid will be transformed into an automation pillar at which TSN serves as the (sole) connectivity provider for control loops which will span the whole automation network from virtualized (/centralized) controllers and the field level. In his talk, Tobias Heer provided an overview of the changes that TSN brings with regard to medium access methods to enable real-time capabilities in Ethernet. While TSN brings significant improvements to wired settings, the subsequent discussion round revolved around the difficulties in achieving this in wireless scenarios. Besides the apparent issues of jamming and/or other attack vectors in wireless control systems, the possibility of trading reliability against capacity and the resulting implications on control algorithms was identified as a research issue.

Conclusion

Throughout the presentations and especially the discussions both during the plenary sessions as well as during off-hour activities, the seminar successfully brought together researchers from control and communication from both academia and industry, and undoubtedly fostered a deeper understanding of the intricate interplay of the disciplines in the research area of CPS/CPN. A variety of open problems and promising research areas were identified, with some in dire need of increased cooperation between the involved fields. This underlines the need in CPS/CPN research for formats valuing open and honest discussions, and both the organizers and the participants hope to be able to continue these discussions in the following years through additional summits and – once the insights gained in this first edition have shown visible impact on the scientific community – possibly another Dagstuhl Seminar. As a concrete follow-up, the organizers and participant James Gross are planning to conduct a seminar in Stockholm/Sweden in 2020 on this diverse research area.

6.35 Empirical Evaluation of Secure Development Processes

Organizers: Adam Shostack, Matthew Smith, Sam Weber, and Mary Ellen Zurko

Seminar No. 19231

Date: June 2–7, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.6.1

Participants: Florian Alt, Adam J. Aviv, Eric Bodden, Michael Coblenz, Tamara Denning, Serge Egelman, Sascha Fahl, Shamal Faily, Tobias Fiebig, Joseph Hallett, Trent Jaeger, Mike Lake, Carl E. Landwehr, Steven B. Lipner, Luigi Lo Iacono, Fabio Massacci, Michelle Mazurek, Brendan Murphy, Brad A. Myers, Xinming (Simon) Ou, Olgierd Pieczul, Heather Richter Lipford, Riccardo Scandariato, Reinhard Schwarz, Adam Shostack, Laurens Sion, Matthew Smith, Walter F. Tichy, Daniel Votipka, Sam Weber, Charles Weir, Laurie Williams, Mary Ellen Zurko

The problem of how to design and build secure systems has been long-standing. For example, as early as 1978 Bisbey and Hollingworth [6] complained that there was no method of determining what an appropriate level of security for a system actually was. In the early years various design principles, architectures and methodologies were proposed: in 1972 Anderson [5] described the "reference monitor" concept, in 1974 Saltzer [7] described the "Principle of least privilege", and in 1985 the US Department of Defense issued the Trusted Computer System Evaluation Criteria [8].

Since then, although much progress has been made in software engineering, cybersecurity and industrial practices, much of the fundamental scientific foundations have not been addressed – there is little empirical data to quantify the effects that these principles, architectures and methodologies have on the resulting systems.

This situation leaves developers and industry in a rather undesirable situation. The lack of this data makes it difficult for organizations to effectively choose practices that will cost-effectively reduce security vulnerabilities in a given system and help development teams achieve their security objectives. There has been much work creating security development lifecycles, such as the Building Security In Maturity Model [1], Microsoft Security Development LifeCycle [3] OWASP [4] and ISECOM [2] and these incorporate a long series of recommended practices on requirements analysis, architectural threat analysis, and hostile code review. It is agreed that these efforts are, in fact, beneficial. However, without answers as to why they are beneficial, and how much, it is extremely difficult for organizations to rationally improve these processes, or to evaluate the cost-effectiveness of any specific technique.

The ultimate goal of this seminar was to create a community for empirical science in software engineering for secure systems. This is particularly important in this nascent of research in this domain stage since there is no venue in which researchers meet and exchange. Currently single pieces of work are published at a wide variety of venues such as IEEE S&P, IEEE EuroS&P, ACM CCS, USENIX Security, SOUPS, SIGCHI, ICSE, USEC, EuroUSEC, and many more. The idea was that bringing together all researchers working separately and creating an active exchange will greatly benefit the community.

Naturally, community-building is a long-term activity – we can initiate it at a Dagstuhl seminar, but it will require continuous activity. Our more immediate goals were to develop a manifesto for the community elucidating the need for research in this area, and to provide *actionable* and *concrete* guidance on how to overcome the obstacles that have hindered progress.

One aspect of this was information gathering on how to conduct academic research which is able to be transitioned and consumed by developers. We felt that all too frequently developer needs aren't fully understood by academics, and that developers underestimate the relevance of academic results. Our information gathering will help foster mutual understanding between these two groups and we specifically looked for ways to build bridges between them.

A second obstacle which we aimed to address is how to produce sufficiently convincing empirical research at a foundational level as well as in the specific application areas. Currently there is no consensus on what are ecologically valid studies and there are sporadic debates on the merits of the different approaches. This seminar included a direct and focused exchange of experience and facilitated the creation of much needed guidelines for researchers. In accordance with our bridge building, we also looked at what developers find convincing, and how that aligns with research requirements.

Seminar Format

Our seminar brought together thirty-three participants from industry, government and both the security and software engineering academic communities. Before the seminar started we provided participants with the opportunity to share background readings amongst themselves.

We began our seminar with level-setting and foundational talks from industrial, software engineering and security participants aimed to foster a common level of understanding of the differing perspectives of the various communities.

Following this the seminar was very dynamic: during each session we broke into break-out groups whose topics were dynamically generated by the participants. The general mandate for each group was to tackle an aspect of the general problem and be actionable and concrete: we wished to avoid vague discussions of the difficulties involved with studying secure development but instead focus on how to improve our understanding and knowledge. After each session we met again as a group and summarized each group's progress.

At the conclusion of the seminar we brought together all the participants in a general discussion about further activities. In all, a total of eighteen further activities, ranging from papers to research guideline documents, were proposed and organized by the participants.

References

1 Building security in maturity model. http://www.bsimm.com/.

2 Isecom. http://www.isecom.org.

3 Microsoft security development lifecycle. https://www.microsoft.com/en-us/securityengineering/sdl/.

4 Owasp. https://www.owasp.org.

5 Anderson, J. P. Computer Security Technology Planning Study, Volume II. Tech. Rep. ESD-TR-73-51, 1972.

6 Bisbey, R., and Hollingworth, D. Protection analysis: Final report. *Information Sciences Institute, University of Southern California: Marina Del Rey, CA, USA, Technical Report ISI/SR-78-13* (1978).

7 Saltzer, J. Protection and the control of information sharing in Multics. *Communications of the ACM 17*, 7 (1974), 388–402.

8 United States Department of Defense. Trusted computer system evaluation criteria (orange book).

6.36 Ubiquitous Computing Education: Why, What, and How

Organizers: Audrey Girouard, Andrew L. Kun, Anne Roudaut, and Orit Shaer

Seminar No. 19232

Date: June 2–7, 2019 | Dagstuhl Seminar
Full report – DOI: 10.3830/DagRep.9.6.26

Participants: Michael Beigl, Andrea Bianchi, Anke Brock, Jessica Cauchard, Ruzanna Chitchyan, Jeremy Cooperstock, Ellen Yi-Luen Do, Audrey Girouard, Eva Hornecker, Miriam Konkel, Andrew Kun, Sylvain Malacria, Nicolai Marquardt, Amanda McLeod, Donald McMillan, Timothy Merritt, Caitlin Mills, Simon Perrault, Thomas Pietrzak, Michal Rinott, Anne Roudaut, Albrecht Schmidt, Oliver Schneider, Orit Shaer, Jakub Sypniewski, Aurélien Tabard, Brygg Ullmer, Vicky Zeamer

This document summarizes the insights gathered during the seminar. We first provide an overview of the motivation for this seminar before presenting an overview of the activities that occurred during these five days. We then provide a series of outputs that we gathered in addition to the list of abstract provided on the website.

1. Motivation

Interactive systems are becoming increasingly complex and diversified, often comprised of multiple interconnected devices, with many different functionalities. They are slowly merging within our everyday objects. Such systems are becoming ubiquitous. Ubiquitous computing, or ubicomp, is a multidisciplinary field of study that explores the design and implementation of such embedded, networked computing systems. Due to the novel aspect of the technologies involved and the multidisciplinary nature of skills needed to design such systems, teaching and training new innovators in this field are not well addressed through traditional programs and instruction. Consequently, it is important to ask several questions about the training and education needed to help students become valuable members and leaders of ubicomp teams. Three central questions about ubiquitous computing education emerge: why, what and how, with the goal of enhancing ubicomp education through interdisciplinary perspectives:

- WHY is training in ubicomp needed? Is it enough to train experts in narrow domains (e.g. those who can create low-power embedded circuits, or those who can make usable applications), and then bring them together in teams that will tackle ubicomp problems? Or do we need specialized training that targets ubicomp in addition to domain expertise? There is broad consensus that we do need specialized training, but often this argument is based on intuition and anecdotal evidence. We approach this question by first asking: what are the grand challenges that we expect our students to tackle in the world (e.g. privacy, sustainability) by inventing and developing ubicomp solutions? Next, we ask: who can better address the challenges: teams of domain experts, or teams where at least some team members have specialized ubicomp education? Answers to these questions will clearly identify problems that might exist with current ubicomp educational approaches.
- WHAT should constitute training in ubicomp? Once we identify the grand challenges, we need to ask further questions. What are the values, knowledge, and skills we should train students in ubicomp? What are the topics that should be covered? How do these depend on the background of students or their degree program? Answers to these types of questions will allow us to set goals for ubicomp education.
- HOW should we teach and engage a diverse body of students? Once we identify specific goals for ubicomp education, we need to ask ourselves how those goals can be achieved. How does the unique nature of ubicomp challenge the current pedagogical approaches? How can we create new pedagogical approaches for teaching and training in ubiquitous computing? Answers to these types of questions will help create the appropriate tools to reach our ubicomp education goals.

2. Overview of the activities

Our goal was to create a community to support new forms of teaching, training, and learning in ubiquitous computing. Our activities were centered on our main questions:

- Day 1, we explored the WHO and WHY. Each participant presented briefly their research and current teaching, and highlighted what they see are the main challenges for teaching

ubicomp in the morning. We then brainstormed and discussed why is is important to rethink the way we teach ubicomp material and what are the grand challenges associated to this change.

- Day 2, we explored the WHAT. In groups, we defined the curriculum for Ubicomp education for different types of students, different degree levels, as well as identified what are the learning goals. One discussion that came up relating to the limits of Ubicomp material, specifically how complex it currently is to define what is ubicomp.
- Day 3, we explored the HOW, and particularly brainstormed about the challenges related to ubiquitous education. Participants generated a list of their current active learning methods or tools and exchanged them in a speed dating fashion with each other.
- Day 4, we explored further the HOW. In groups, we developed and experienced new active learning pedagogies on ubiquitous computing topics. We also discussed pedagogies for academic ubicomp programs.
- Day 5, we wrapped up the seminar and plan for concrete actions for the future, in particular, ideas for the next Dagstuhl seminar.

3. The challenges of teaching Ubicomp (WHY)

Figure 6.5 illustrates the grand challenges of teaching Ubicomp from a motivation point of view. We have identified several themes including (1) who is the audience in terms of diversity, motivation, population; and how (2) these aspects particularly impact their engagement and what methods can we use to better engage with students. We also talked about the difficulty that Ubicomp brings in terms of being a multi-disciplinary field and we highlight the fact that it is difficult to choose (3) which topics should be covered and which ones should not be covered in a particular case. What are the boundaries of Ubicomp? In fact, our discussions highlighted that there is not a clear (4) definition of Ubicomp. We talked about (5) issues with the high workload of both teaching and learning about ubicomp, and how research-led teaching could alleviate some of these issues. We discussed (6) scale issues, i.e. how to teach to a large number of students (and provide feedback) when it seems that certain aspects of Ubicomp teaching (e.g. workshop activities) can only be taught to smaller groups. We pointed out the issues of (7) space and that Ubicomp teaching is based on traditional classroom but also new types of spaces such as workshops, hackerspaces, and maker spaces. Furthermore, we discussed other media types such as (8) online lectures. We also discussed more general topics such as (9) the impact of ubicomp (e.g. on business and industry) and the future of universities and how this relates to ubicomp education.

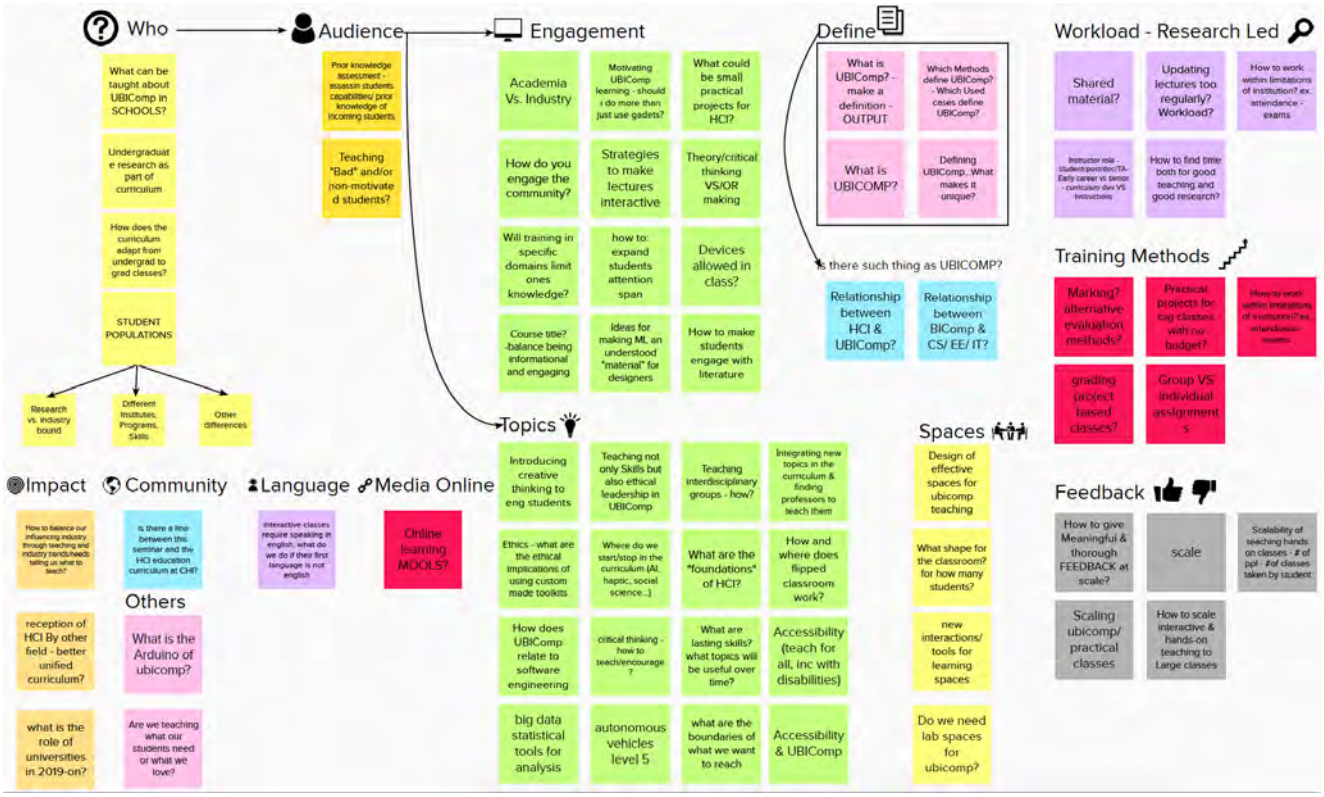

Fig. 6.5
Grand challenges of teaching Ubicomp (from a motivation point of view).

4. The Ubicomp curriculum (WHAT)

Table 6.1 illustrates the topics central to Ubicomp Education brainstormed during the seminar. We split the participants (including the organizers) in four groups designing curriculum (standalone lecture or program) for different students (undergraduate UG or postgraduate PG) and technical (Computer Science) or non-technical (Interdisciplinary) background). We wish for this document (that we also plan to put onto our online web platform) to be used as guidelines for teachers in order to provide a better and unified Ubicomp curriculum across different institutions and countries.

5. Existing active learning methods for Ubicomp (HOW)

Figure 6.6 illustrates the grand challenges of teaching Ubicomp from a method's point of view. This was the result of a brainstorming with participants following the curriculum creation. We found that (1) managing the workload was a theme recurring again (as we also mentioned it in the initial brainstorming in Figure 6.5). We noted that one difficulty of teaching Ubicomp was (2) the lack of differentiation with other CS, HCI or Design teaching material. We also though this could create issue in (3) attacking certain types of students and that possibly, depending on the demographic, different terminology (Ubicomp, Interactive Systems, Interaction Design etc.) might be used. We raised issues in (4) evaluation and assessment potentially raised by (5) the interdisciplinarity of the community which makes it hard to assess student but also to teach so diverse material. We discussed issue in (6) engaging with students and enforcing skill acquisition (surface vs. deep learning). Finally we also add other issues such as (7) scaling of students, (8) project styles, (9) reaching to real end-users, (10) having input from industry and the (11) format of the lecture (e.g. online). We finally discussed about the issues raised by (12) admin and physical resources.

To build on participants' past and current experience regarding education, we also asked them to share both memorable experiences as well as active learning methods. For the former, we wanted to gather memorable educational moments, anecdotes that stayed with participants long after, as means to both remember the impact that we have on others, as well as get inspiration when designing new activities or methods.

We finally asked participants to share three teaching active learning exercises or methods, ones they currently use in their teaching materials, or ones they experienced in the past. They shared their methods, in a one on one, speed-networking format. In two minutes, they explained one of their ideas to another participant. After the speed-dating, participants placed their basic descriptions on a board and voted for the ones that seemed relevant to their courses. This activity sparked interest for material sharing among participants and ideas on how everyone could implement some of the approaches in their own contexts.

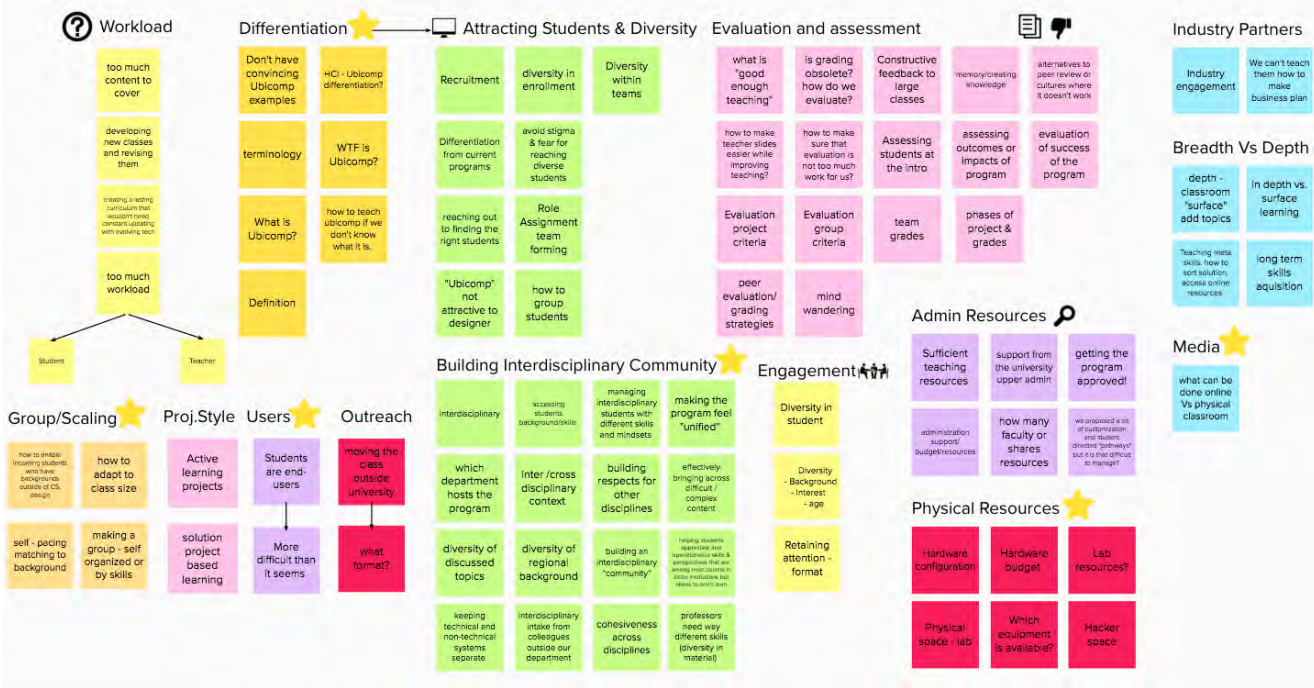

Fig. 6.6
Grand challenges of teaching Ubicomp (from a method point of view). A star corresponds to a challenge specific to Ubicomp while other are general to teaching.

6. Innovative active learning methods for Ubicomp (HOW)

The next main activity focused on generating new educational material that may be difficult to generate, or missing, from a current curriculum. Participants formed six groups, they selected a topic, and investigated new active learning methods as well as initial teaching material related to the topic. The specific topics for each group were selected from topics and challenges highlighted earlier in the seminar. Next, groups formed pairs of groups, and each group tested their content and methods on other group and received feedback, before iterating on their design. Finally, teams

Table 6.1
Ubicomp Curriculum by themes, topics and types of students.

Theme	Topic	Course UG Interdisciplinary	Course UG Technical	Course PG Interdisciplinary	Course PG Technical	Course PG *Architecture*	Program PG Technical	Program UG Interdisciplinary
Introduction	Visions of ubicomp	x	x	x	x	x	x	x
	Ubicomp Fundamentals	x	x	x	x	x	x	
	History of computing	x	x		x	condensed		x
	Interfaces	x	x	x		x		
Tech	Hardware skills	basic	x	x (prior knowledge, class, or bootcamp)	x	basic	x (prior knowledge, class, or bootcamp)	x
	Software skills	basic	pre requisite	x (prior knowledge, class, or bootcamp)	x	basic	x (prior knowledge, class, or bootcamp)	x
	Fabrication techniques	basic	x		x	basic since prior knowledge	x	
	Interaction techniques, modalities	basic	x		x	overview	x	
	Electronics	x	x	x	x			
	Sensors	x	x	x				
	Actuators	x	x	x				
	Location tracking		x	x				
	IoT		x	x			x	
	Signal processing				x	use toolkits	elective	
	Network	basic	x	x			x	x
	Activity recognition				x			
	Communication protocols				x			
	System building		x		x			
	Infrastructures				x			
	Appliance design		elective		x			
	Displays						elective	
	Content Awareness			x				
	Testing, certification, ISO				x		x	
Methods / Design	User Centered Design (UCD)		x	possibly pre-requisite	x			
	Sketching + design	x	x	x				x
	Prototying Methods	x	x	x	x			x
	Evaluation		x	x	x	x	x	x
	Inclusive/accessible design			elective	elective		elective	elective
	Statistics				pre-requisite		x	x
Specific domains	AI/ML	basic	optional	optional			elective	
	Human Augmentation	basic	optional	optional				
	Data Science/Analytics			optional	x		elective	
	Robotics						elective	
	Sustainability					x		x
	Autonomous systems		advanced					
	Entrepreneurship			elective				elective
Implications	Ethical considerations	x	x	x	pre-requisite	x	x	x
	Security	x	x	x	pre-requisite	x	x	x
	Privacy	x	x	x	pre-requisite	x	x	x
Communication	Demonstration of product		x	x	x	x	x	x
	Writing/describing product		x	x	x	x	x	x

presented a summary of their new materials to the group (Table 6.2).

From the discussions following up the presentation we also noted some actions to do:

- Create a repository or playlist (youtube) or videos that can be used within the community and define what is Ubicomp.
- Ask participants to upload a 2 minutes video of their definition of Ubicomp that can be used in class to show the variety of what people think is Ubicomp.

7. Future Steps

Although this seminar addressed many questions the organisers had originally highlighted, it also opened new exciting directions to explore and new challenges. To start addressing them we identified the main following avenues for future work and future events:

- Follow-up Dagstuhl seminar on writing a textbook
- Follow-up Dagstuhl seminar focussing on the industrial side, e.g. what skills do students need for the society we will built in 5/10/20/50 years?
- Using the website to keep the community alive as well as the access to material, and also create a video channel to create a repository of ubicom examples.
- We also have discussed about 3-4 follow up papers to be written among participants and organizers.

8. Reading list

We collected a reading list that addresses the why, what, and how of ubicomp education, designed for educators.

- Pervasive Computing Education, Audrey Girouard, Andrew L. Kun, Anne Roudaut, Orit Shaer, and Andrew L. Kun, IEEE Pervasive Computing, Oct, 2018.
- Teaching Pervasive Computing in Liberal Arts Colleges, Orit Shaer, and Evan M. Peck IEEE Pervasive Computing, Jul, 2018.
- The Fuzzy and the Techie: Why the Liberal Arts Will Rule the Digital World, Scott Hartley
- Fixing Tech's Ethics Problem Starts in the Classroom, Stephanie Wykstra, The Nation.
- The Pervasive, Embedded, and Mobile Computing Curriculum – Preparing Computer Science Students for the Technology of the Future, Jakob Bardram, 2012
- Bridging the Gap Between Teaching and Research: A Case Study for Engineering & Applied Science, Anne Roudaut, Higher Education Pedagogies 2019
- Krumm, J. (Ed.) Ubiquitous computing fundamentals. CRC Press, 2010 (fairly outdated by now)
- Rowland et al. Designing Connected Products. O'Reilly 2015 (on Design of IoT products, with a broad range of topics ranging from networking aspects, architecture to product design)
- Landay, J. A., & Borriello, G. (2003). Design patterns for ubiquitous computing. *Computer*, *36*(8), 93-95.
- Electronics books from "Make" .
- The radar diagram https://scottwhyoung.com/teaching/information-ethics-privacy-spring- 2018/

Table 6.2
Innovative content and activities for Ubicomp education.

Topic	Content	Activities
What is Ubicomp (Ellen, Jakub, Michael)	▪ Multi-media video presenting Ubicomp and that the whole community can use + show examples of real life Ubicomp (or fictional like on movie) ▪ Have a good and bad example of Ubicomp repository video ▪ They showed a video of the construction worker in traditional set up and new (with new digital tool) ▪ Introduce key concepts (e.g. calm technology, Weiser, disappearing computer etc.) ▪ ILO: identify concepts in showed example / explain the example using appropriate terminology ▪ What is Ubicomp " let's people define HCI. It is too big to propose a definition because the edges are so blurry. But this is may be not a bad thing (although students may ask us what it is). ▪ What is ubicomp? May be it is about designing interconnected devices while removing the stress of having so many devices to deal with (from a user point of view)	▪ Make them watch videos and read paper and write a definition of Ubicom (plus sketch) ▪ Definition of terminology (e.g. Cliff notes or cheatsheet) ▪ Give examples of categories of things that can disappear (e.g. task, physical computer) and ask students how would you make it disappear. ▪ Repository of video or youtube playlist?
Electronics (Tim, Brygg, Thomas)	▪ Use three examples ("schools") of building new Ubicomp devices using simple electronics and use these examples all along the course. ▪ Three layers of increasing complexity to focus on are sensing, actuating, computing. ▪ Having a 30 minutes version of this as well.	
Electronic (Nicolai, Jessica, Jeremy, Andrea, Simon	▪ 110 minute lecture about sensors part of a more general course on physical computing ▪ Here focus is on hardware so code is not the focus and is prepared (download and pre-assemble circuit). ▪ Basic of resistors (sensing as a variation of resistance)	▪ Before class: they get material and they assemble and glue together a sensors (plus papers material). They will end up with different sensors ▪ Motivating by showing famous consumer devices and show how it is inside and this is "simple" ▪ Show what's inside a resistor ▪ Find out about the mapping between the sensing value we are getting and the actual response (e.g. acrylic overlays to change the light coming to a sensor) ▪ Collect all broken device and identify things inside.
Interfaces (Anke, Michal, Sylvain, Donald)	▪ Design space to describe the dimension of interacting with Ubicomp. Show example of Ubicomp visions and show where they are in the design space. ▪ It may be good to show old video trying to predict the future to show what things are changing, may be why defining Ubicomp is hard. Are they future of the past?	▪ Activity 1 "understanding possible ubicomp interfaces: gave existing material and let students place them in the design space" (the mock up activity show this may be a good activity). ▪ Activity 2 "Desining with ghosts" (Donald use it already). This is kind of a wizard of oz design where two students are ghosts and actuated things around and the user is blindfolded. The ghost help them through the world by actuating things around. ▪ Bank of videos example
Research methods (Caitlin, Miriam, Vicky)	▪ For interdisciplinary class (45 min) ▪ ILO: identifying research question, making a hypothesis, identifying variables, operational definition and study design.	▪ Give hypothesis example, .e.g "the time of the day will improve your grade" . Ask "Is your hypothesis falsifiable" ? ▪ Activity around what is measured in terms of the dependant variable and how to make sure the measurement fits with the hypothesis. ▪ Watch a short video, come up with a question "does self driving cars cause more accidents" . ▪ Giving example of wrong design or wrong set up and make them find why and where are the problems.
Implication (Aurelien, Eva, Ruzanna)	▪ Use integrated approach to talk about the implication during the entire course rather than in a single lecture ▪ Topics as dimensions to discuss within a diagram they proposed (see Radar diagram in the reading list) ▪ Show the "what our smart devices know about you" (TED talk)	▪ Use case studies to discuss impact of technology or architecture choices on impact topics (from the diagram) as well as other factors. ▪ Exercise with critical review of certain system can be done. ▪ Repository of use case examples (e.g. Amazon dash video, google glass, autonomous cars, Alexa etc.) ▪ "Common thread" with a visual so students know when the implication is discussed during an entire program.

6.37 25 Years of the Burrows-Wheeler Transform

Organizers: Travis Gagie, Giovanni Manzini, Gonzalo Navarro, and Jens Stoye

Seminar No. 19241

Date: June 10–14, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.6.55

Participants: Jarno Alanko, Hideo Bannai, Paola Bonizzoni, Christina Boucher, Marilia Braga, Anthony J. Cox, Fabio Cunial, Jackie Daykin, Richard Durbin, Gabriele Fici, Johannes Fischer, Travis Gagie, Pawel Gawrychowski, Simon Gog, Roberto Grossi, Wing-Kai Hon, Tomohiro I, Juha Kärkkäinen, Dominik Kempa, Tomasz Kociumaka, Dominik Köppl, Ben Langmead, Zsuzsanna Liptak, Veli Mäkinen, Sabrina Mantaci, Giovanni Manzini, Ian Munro, Gene Myers, Gonzalo Navarro, Yakov Nekrich, Enno Ohlebusch, Kunsoo Park, Nicola Prezza, Knut Reinert, Giovanna Rosone, Kunihiko Sadakane, Leena Salmela, Marinella Sciortino, Rahul Shah, Sandip Sinha, Jouni Sirén, Tatiana Starikovskaya, Jens Stoye, Sharma V. Thankachan, Rossano Venturini

Dagstuhl Seminar 19241 marked the 25th anniversary of the publication of the Burrows-Wheeler Transform (BWT), which has had a huge impact on the fields of data compression, combinatorics on words, compact data structures, and bioinformatics. The 10th anniversary in 2004 was marked by a workshop at the DIMACS Center at Rutgers (http://archive.dimacs.rutgers.edu/Workshops/BWT) organized by Paolo Ferragina, Giovanni and S. Muthukrishnan, and it is exciting to see how far we have come. In the past 15 years, interest in the BWT has shifted from data compression to compact data structures and bioinformatics, particularly indexing for DNA read alignment, but seven of the 33 participants of that workshop (including Giovanni) also attended this seminar. Unfortunately, Professor Gørtz fell ill at the last minute and emailed us on June 11th to say she couldn't attend, but everyone else on the final list of invitees was present for at least some of the seminar (although not everyone made it into the photo). In total there were 45 people (listed at the end of this report) from 13 countries, including ten women, six junior researchers and two researchers from industry. By happy coincidence, the seminar started the day after Gonzalo's 50th birthday, so we were able to celebrate that as well. We thank Professor Sadakane for the photos shown in Figures 6.7 and 6.8.

The schedule, shown in Figure 6.9, featured an introduction, 12 talks, three panel sessions and a closing. The talks were all timely and reflected the active and dynamic research being carried out on the BWT:

- Giovanni's introduction was a more in-depth version of his invited talk from DCC '19;
- Veli Mäkinen surveyed pan-genomic indexing, including work published in *BMC Genomics* last year;
- Richard Durbin surveyed results based on the positional BWT, published in *Bioinformatics* in 2014;
- Jouni Sirén presented work included in a *Nature Biotechnology* article last year;
- Christina Boucher surveyed compact data structures for de Bruijn graphs, including work from an ISMB/ECCB 2019 paper;
- Gonzalo Navarro reviewed BWT-based indexes, including work from a SODA '18 paper;
- Sandip Sinha presented work from a STOC '19 paper;
- Dominik Kempa presented work from another STOC '19 paper;
- Sharma Thankachan presented work from an ESA '19 paper;
- Nicola Prezza presented work from a STOC '18 paper;
- Marinella Sciortino gave a version of her invited lecture for IWOCA '19 a month later;
- Giovanna Rosone presented results about two extensions of the BWT, including work from a WABI '18 paper, now published in *Algorithms for Molecular Biology*;
- Dominik Köppl presented work from a CPM '19 paper.

We later received all the abstracts but one.

Fig. 6.7
Gonzalo with his birthday cake (featuring a BWT).

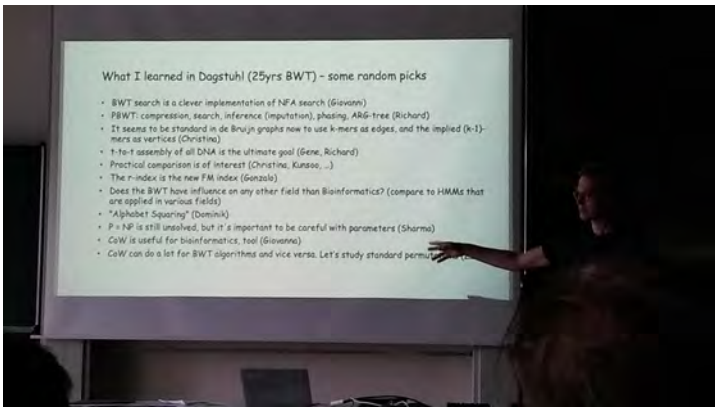

Fig. 6.8
Jens reviewing some points raised during the seminar.

	MONDAY	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	
07:30		BREAKFAST	BREAKFAST	BREAKFAST	BREAKFAST	
09:00		INTRO	ALG TALK 1	CoW TALK 1	WORK...	
09:45		BIO TALK 1	ALG TALK 2	CoW TALK 2		
10:30		BIO TALK 2	ALG TALK 3	CoW TALK 3		
11:15		BIO TALK 3	ALG TALK 4	CoW TALK 4		
12:15		LUNCH	LUNCH	LUNCH	LUNCH	
13:45		BIO TALK 4				
14:00			ALG PANEL	CoW PANEL		
14:30		BIO PANEL				
15:00			WORK!	CLOSING		
15:30	CAKE	CAKE	CAKE	CAKE		
16:00	WORK?	WORK	WORK!!	WORK!!!		
18:00	DINNER (buffet)	DINNER	DINNER	DINNER		
20:00	CHEESE?	CHEESE	CHEESE	CHEESE		

INTRO	Giovanni	BIO PANEL	ALG PANEL	CoW PANEL
BIO TALK 1	Veli	(Pan-genomic) alignment	Ben	Ian	Gabriele
BIO TALK 2	Richard	PBWT	Gene	Inge (chair)	Hideo
BIO TALK 3	Jouni	GBWT	Knut	Johannes	Jackie
BIO TALK 4	Christina	de Bruijn graphs	Kunsoo	Rahul	Pawel
ALG TALK 1	Gonzalo	r-index	Paola	Roberto	Sabrina (chair)
ALG TALK 2	Sandip	Local decodability	Richard	Simon G	Tomasz
ALG TALK 3	Dominik	BWT construction	Tony (chair)		Zsuzsa
ALG TALK 4	Sharma	Wheeler graphs			
CoW TALK 1	Nicola	String attractors	Jens chairs BIO talks		
CoW TALK 2	Marinella	Combinatorial properties	Giovanni chairs ALG talks		
CoW TALK 3	Giovanna	eBWT / BWT similarity	Travis chairs CoW talks		
CoW TALK 4	Dominik	Bijective BWT			
CLOSING	Jens				

Fig. 6.9
The original seminar schedule. Inge Li Gørtz was unable to attend and so Tatiana Starikovskaya chaired the *Algorithms and Data Structures* panel. The talks and panel on *Bioinformatics* were held on the first day and those on *Algorithms and Data Structures* on the second day to accommodate participants' schedules.

6.38 Distributed Computing with Permissioned Blockchains and Databases

Organizers: C. Mohan, Beng Chin Ooi, and Gottfried Vossen

Seminar No. 19261

Date: June 23–28, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.6.69

Participants: Divyakant Agrawal, Alysson Neves Bessani, Jeeta Ann Chacko, Lei Chen, Mariano P. Consens, Tien Tuan Anh Dinh, Alan Fekete, Michael J. Franklin, Vincent Gramoli, Krishna P. Gummadi, Michael Huth, Hans-Arno Jacobsen, Murat Kantarcioglu, Srinivasan Keshav, Shahan Khatchadourian, Dilip Krishnaswamy, Juho Lindman, Eric Lo, Alexander Löser, Dumitrel Loghin, Bernhard Mitschang, C. Mohan, Hart Montgomery, Pezhman Nasirifard, Beng Chin Ooi, Torben Bach Pedersen, Dennis Przytarski, PingCheng Ruan, Gabriela Ruberg, Mohammad Sadoghi Hamedani, Yong Tang, Gottfried Vossen, Li Xiong, Feida Zhu

The topic of blockchains, and in particular that of permissioned blockchains, has rapidly gained interest in both the industrial and the research communities in recent years. It particularly pertains to situations where trust among several parties that are about to do business together is difficult to establish (e.g., due to organizational, financial, or timing reasons) or impossible to establish at all. A blockchain is a decentralized, distributed ledger that consists of immutable blocks containing transactions that can be accessed by any party, and that provides trust via replication over all nodes and an agreed-upon execution order of the transactions. Of particular interest are permissioned blockchains in which the associated parties are known and authenticated, yet still do not fully trust each other.

Many applications have shown interest in the concept of blockchains, since the situation just described applies to many real-world scenarios, including (global) supply chains, the Internet of Things, connected cars, manufacturing, banking, and healthcare. As a consequence, a number of players in the IT industry work on a development of the technology, and several consortia have been formed to advance the technology across industries, among them Hyperledger and R3. Moreover, a number of companies have released Blockchain-as-a-Service (BaaS) platforms, including IBM, Oracle, Amazon, Baidu, and Alibaba.

The technology has many links into the database community; however, the situation is basically like it was in the database area many years ago, when only a few systems had been released but users were on their own to figure out how to use them effectively. As the seminar has shown, many interesting issues remain to be solved, and there is a wide variety of aspects and research issues currently under investigation. Of these, the following were discussed:

- Blockchain scalability w.r.t. transaction throughput, one of the main roadblocks to business adoption
- Transaction ordering and endorsement, consensus of transaction commit
- Adjustments to the Proof of Work (PoW) consensus mechanism, other optimizations to consensus algorithms (e.g., Byzantine consensus) in the presence of transaction failures and in light of scalability
- Block validation
- Languages for smart-contract specification (e.g., Sandcastle SQL and Solidity)
- Amendments to Hyperledger Fabric, such as channels
- Cross-chain swaps using hashed timelocks
- Energy efficiency of blockchain applications

In addition, several participants reported on various working applications of blockchain technology.

6.39 Astrographics: Interactive Data-Driven Journeys through Space

Organizers: Alyssa A. Goodman, Charles D. Hansen, Daniel Weiskopf, and Anders Ynnerman

Seminar No. 19262

Date: June 23–26, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.6.95

Participants: Brian Abbott, Hidehiko Agata, Emil Axelsson, Thomas Boch, Alexander Bock, Dave Brown, Melvyn Davies, Carter Emmart, Jackie Faherty, Andreas Gerndt, Alyssa A. Goodman, Charles D. Hansen, Tom Kwasnitschka, David H. Laidlaw, Marcus A. Magnor, Thomas Müller, Joshua Eli Goldston Peek, Lucian Plesea, Sebastian Ratzenböck, Thomas P. Robitaille, Filip Sadlo, Wolfgang Steffen, Gabriel Stöckle, Mark Subbarao, Edwin A. Valentijn, Daniel Weiskopf, Ryan Wyatt, Anders Ynnerman

For the majority of human existence, the visual language has been successfully used to communicate complex ideas that span across borders of knowledge, experience, age, gender, culture, and time. These aspects also make it an effective form of expressing workflows in scientific data analysis as well as the communication of scientific discoveries to broad audiences. The Dagstuhl Seminar 19262 brought together researchers from computer science, content producers, learning and communication experts, and domain experts from astronomy and astrophysics to define the emerging field of interactive visualization of space exploration and astronomy, referred to as Astrographics. This seminar played an important role in the ongoing process of removing the clear division between using visualization to enable scientific discoveries by subject-matter experts (exploratory visualization) and using visual representations to explain and communicate the results of such exploratory science to a greater, general audience (explanatory visualization). Designing the available visualization tools to serve both roles at the same time increases the overlap between these two aspects of visualization and allows scientists to better explain their findings and, at the same time, enables the general public to use similar tools for their own, guided, discovery and actively participate in the scientific process. The field of astronomy and astrophysics has been at the forefront of this process since the beginning as it is a primary example of a domain in which exploratory and explanatory visualizations have served important but distinct roles. For this reason, astrographics was chosen as the domain in which to explore the challenges and opportunities that arise when combining exploratory and explanatory techniques. The bulk of work in this seminar occurred in focussed break-out sessions that reported their findings back to the group and opened up the topics for joint discussions. Topics of these break-out sessions included discussions on better integration of software tools, improvements of analysis tools, preparing astrographics software packages to improve the quality of public presentations, the ability of sharing presentations both in spatially distant locations as well as saving them for later playback. Finally, there was a working group to work on a decadal white paper for astronomy [1].

References

1 Jacqueline K. Faherty, Mark SubbaRao, Ryan Wyatt, Anders Ynnerman, Neil deGrasse Tyson, Aaron Geller, Maria Weber, Philip Rosenfield, Wolfgang Steffen, Gabriel Stoeckle, Daniel Weiskopf, Marcus Magnor, Peter K. G. Williams, Brian Abbott, Lucia Marchetti, Thomas Jarrett, Jonathan Fay, Joshua Peek, Or Graur, Patrick Durrell, Derek Homeier, Heather Preston, Thomas Müller, Johanna M Vos, David Brown, Paige Giorla Godfrey, Emily Rice, Daniella Bardalez Gagliuffi, Alexander Bock. IDEAS: Immersive Dome Experiences for Accelerating Science. arXiv preprint arXiv:1907.05383, 2019

6.40 Graph Colouring: from Structure to Algorithms

Organizers: Maria Chudnovsky, Daniel Paulusma, and Oliver Schaudt

Seminar No. 19271

Date: June 30–July 5, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.6.125

Participants: Isolde Adler, Marthe Bonamy, Nicolas Bousquet, Christoph Brause, Kathie Cameron, Maria Chudnovsky, Konrad Dabrowski, Cemil Dibek, Francois Dross, Esther Galby, Petr A. Golovach, Chinh T. Hoàng, Shenwei Huang, Bart Jansen, Matthew Johnson, Mamadou Moustapha Kanté, Tereza Klimosova, Stefan Kratsch, O-joung Kwon, Bernard Lidicky, Anita Liebenau, Paloma Lima, Daniel Lokshtanov, Peter Maceli, Tomáš Masařík, Sang-il Oum, Daniel Paulusma, Irena Penev, Marcin Pilipczuk, Michal Pilipczuk, Bernard Ries, Paweł Rzążewski, Oliver Schaudt, Ingo Schiermeyer, Pascal Schweitzer, Paul Seymour, Sophie Spirkl, Juraj Stacho, Maya Jakobine Stein, Stéphan Thomassé, Nicolas Trotignon, Zsolt Tuza, Erik Jan van Leeuwen, Gerhard J. Woeginger, Victor Zamaraev

The Graph Colouring problem is to label the vertices of a graph with the smallest possible number of colours in such a way that no two neighbouring vertices are identically coloured. Graph Colouring has been extensively studied in Computer Science and Mathematics due to its many application areas crossing disciplinary boundaries. Well-known applications of Graph Colouring include map colouring, job or timetable scheduling, register allocation, colliding data or traffic streams, frequency assignment and pattern matching. However, Graph Colouring is known to be computationally hard even if the number of available colours is limited to 3.

The central research aim of our seminar was to increase our understanding of the computational complexity of the Graph Colouring problem and related NP-complete colouring problems, such as Precolouring Extension, List Colouring and H-Colouring. The approach followed at the seminar for achieving this aim was to restrict the input of a colouring problem to some special graph class and to determine wether such a restriction could make the problem tractable.

As input restriction, the main focus was to consider hereditary graph classes, which are those classes of graphs that are closed under vertex deletion. Hereditary graph classes provide a unified framework for a large collection of well-known graph classes. The reason for this is that a graph class is hereditary if and only if it can be characterized by a (unique) set $\mathcal{H}$ of minimal forbidden induced subgraphs. This property enables a *systematic* study into the computational complexity of a graph problem under input restrictions. For instance, one can first restrict the input to some hereditary graph class for which $\mathcal{H}$ is small, say $\mathcal{H}$ has size 1 or 2, or for which $\mathcal{H}$ consists of small graphs only.

In line with the seminar's research aim, the seminar brought together researchers from Discrete Mathematics, working in structural graph theory, and researchers from Theoretical Computer Science, working in algorithmic graph theory. In total, 45 participants participated from 14 different countries.

The scientific program of the seminar consisted of 23 sessions: 4 one-hour survey talks, 17 contributed talks of at most thirty minutes and 2 open problem sessions. This left ample time for discussions and problem solving.

Each of the four survey talks covered a particular structural or algorithmic key aspect of the seminar to enable collaborations of researchers with different backgrounds. On Monday, Sophie Sprikl presented a state-of–the-art summary of the Graph Colouring problem for H-free graphs and gave the main ideas and techniques behind an important, recent result in the area, namely a polynomial-time algorithm for colouring P_6-free graphs with at most four colours. On Tuesday, Marcin Pilipczuk gave a tutorial on the framework of minimal chordal completions and potential maximal cliques. This technique plays a crucial role for solving the Maximum Independent Set problem on some hereditary graph classes, but has a much wider applicability. On Wednesday, Bart Jansen gave a presentation on the parameterized complexity of the Graph Colouring problem and related colouring problems. Due to a large variety of possible parameterizatons, Jansen's talk covered a wide range of open problems. On Thursday, Konrad Dabrowski gave an introduction to the clique-width of hereditary graph classes. If a graph class has bounded clique-width, then Graph Colouring and many other NP-hard problems become polynomial-time solvable. Hence, as a first step in the design of a polynomial-time algorithm, one may first want to verify if the clique-width (or any equivalent width parameter) of the graph class under consideration is bounded.

The two general open problem sessions took place on Monday and Tuesday afternoon. Details of the presented problems can be found in the report, together with abstracts of all the talks.

6.41 Real VR – Importing the Real World into Immersive VR and Optimizing the Perceptual Experience of Head-Mounted Displays

Organizers: Marcus A. Magnor and Alexander Sorkine-Hornung

Seminar No. 19272

Date: June 30– July 3, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.6.143

Participants: Philippe Bekaert, Tobias Bertel, Brian Cabral, Susana Castillo Alejandre, Darren Cosker, Douglas Cunningham, Peter Eisert, Atanas Gotechev, Adrian Hilton, Moritz Kappel, Hansung Kim, Christian Lipski, Marcus A. Magnor, Anjul Patney, Johanna Pirker, Christian Richardt, Dieter Schmalstieg, Alexander Sorkine-Hornung, Frank Steinicke, Qi Sun, Christian Theobalt, James Tompkin, Marco Volino, Li-Yi Wei, Gordon Wetzstein, Erroll Wood, Feng Xu

The Dagstuhl seminar brought together 27 researchers and practitioners from academia and industry to discuss the state-of-the-art, current challenges, as well as promising future research directions in Real VR. Real VR, as defined by the seminar participants, pursues two overarching goals: facilitating the import of real-world scenes into head-mounted displays (HMDs), and attaining perceptual realism in HMDs. The vision of Real VR is enabling to experience movies, concerts, even live sports events in HMDs with the sense of immersion of really "being-there", unattainable by today's technologies.

In the welcome and overview session, the participants collectively decided on the seminar program for the following days. In total, the seminar program included the overview session, three research presentation sessions, two breakout sessions including a demo track, two sessions for one-on-one discussions and individual exchange, one session for writing up the results, plus the summary and closing session.

To kick off the seminar, Alexander Sorkine-Hornung from Oculus VR presented the latest developments from an industrial perspective. He gave insights from the development of the just-released Oculus Quest and Oculus Rift S HMDs. In the research presentation sessions, 21 participants gave talks on their work. Participants also met in smaller groups in the breakout sessions to discuss the specific challenges of these fields in more detail. In due course, it became apparent that Real VR concerns research challenges in a number of different fields:

- Capture
- Reconstruction & modeling
- Rendering & perception
- Display technologies
- Interaction & virtual avatars
- Production & applications

Some exemplary results of the seminar on these topics were:

The persistent lack of consumer-market, i.e. affordable, mid- to high-resolution 360-degree video cameras to **capture** dynamic real-world scenes omnidirectionally still hamper research and development in Real VR. So far, research groups largely build their own custom-designed omnidirectional video cameras. Prominent examples include the omnidirectional camera designs by the group of Philippe Bekaert from Hasselt University, Belgium, and the top-of-the-line Manifold camera presented by Brian Cabral from Facebook. Besides novel devices, also simpler recording methods are sought, e.g. by Tobias Bertel and Christian Richardt at Bath, in order to capture real-world content more casually.

On **scene reconstruction and representation**, the jury is still out whether omnidirectional video should be considered to represent sparse light field data with dense depth/disparity as side information, or whether panoramic footage should (and could) be processed to provide full 3D geometry representations of the scene. As pointed out by Atanas Gotchev from TU Tampere, Marco Volino from the University of Surrey, and Christian Richardt from the University of Bath, both forms of representation have their respective advantages and drawbacks, e.g. when aiming to augment the real scene with additional virtual content.. Memory requirements and real-time streaming bandwidth requirements are challenging in either case.

The form of scene representation also determines which **rendering** approaches are viable. For 3D rendering, Dieter Schmalstieg from Graz presented his Shading Atlas Streaming approach to efficiently divide shading and rendering computation between server and client. To make use of visual **perception** characteristics in wide field-of-view HMDs, on the other hand, foveated rendering approaches, e.g. based on hardware ray tracing and accelerated machine learning, as presented by Anjul Patney from NVidia, have great potential. As shown by Qi Sun from Adobe, perceptual methods like saccade-aware rendering can also

be used to enable walking through huge virtual worlds while actually not leaving the confines of one's living room. To render from dense depth-annotated 360-deg video, in contrast, advanced image-based warping methods and hole-filling approaches are needed, as was convincingly outlined by Tobias Bertel from the University of Bath.

Gordon Wetzstein from Stanford University presented how future HMDs will become even more realistic by overcoming current limitations of near-eye **displays**, in particular the vergence-accommodation conflict. Along similar lines, Hansung Kim from the University of Surrey showed how spatial audio enhances perceived VR realism even more.

Social **interaction** in the virtual world requires having digital doubles available. The elaborate steps needed to create convincing human avatars from real-world people were outlined by Feng Xu from Tsinghua University, Darren Cosker from the University of Bath, Christian Theobalt from MPII, and Peter Eisert from TU Berlin, covering the full range of human face, hand and body capture, reconstruction, and modeling. To interact with objects in virtual space, on the other hand, Erroll Wood from Microsoft Cambridge described how hand motion and gestures can be reliably tracked and identified in real-time by the upcoming HoloLens 2 device. Also based on real-time tracking, Li-Yi Wei from Adobe presented a system that enables presenters to augment their live presentation by interacting with the shown content in real-time using mere hand gestures and body postures.

Regarding **content production and applications**, Christian Lipski from Apple presented the ARKit software framework developed for creating captivating augmented reality experiences. James Tompkin from Brown University presented work on multi--view camera editing of Real VR content during post-production. Johanna Pirker from TU Graz showed how virtual reality can be paired with human-computer interaction to enhance learning experiences in the physics classroom. Production aspects and cinematic VR experiences were also considered prominent drivers of contemporary Real VR research by other presenters, e.g. Marco Volino, Darren Cosker, Philippe Bekaert, Peter Eisert and Brian Cabral.

Practically experiencing the new, tetherless Oculus Quest brought along by Alexander Sorkine-Hornung in the **demonstration track** made impressively clear how free, unrestricted user motion extends the usability and acceptance of VR tremendously, made possible by the pass-through view feature of this HMD.

Finally, in the coming months, a number of seminar participants will compile an edited book volume on the state-of-the-art in Real VR that Springer has already agreed to publish as part of their well-known Lecture Notes on Computer Science (LNCS) Survey Series.

6.42 Notional Machines and Programming Language Semantics in Education

Organizers: Mark Guzdial, Shriram Krishnamurthi, Juha Sorva, and Jan Vahrenhold

Seminar No. 19281

Date: July 7–12, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.7.1

Participants: Thomas Ball, Titus Barik, Brett A. Becker, Neil C. C. Brown, Franziska Carstens, Luke Church, Kathryn Cunningham, Paul Denny, Brian Dorn, Benedict du Boulay, Rodrigo Duran, Barbara Ericson, Sally Fincher, Kathi Fisler, Robert L. Goldstone, Mark Guzdial, Reiner Hähnle, Matthias Hauswirth, Arto Hellas, Geoffrey L. Herman, Felienne Hermans, Matthew C. Jadud, Johan Jeuring, Antti-Juhani Kaijanaho, Philipp Kather, A. J. Ko, Shriram Krishnamurthi, Thomas D. LaToza, Colleen Lewis, Elena Machkasova, Craig Miller, Andreas Mühling, Markus Müller-Olm, Greg Nelson, Andrew Petersen, Joseph Gibbs Politz, André L. Santos, Carsten Schulte, Otto Seppälä, R. Benjamin Shapiro, Juha Sorva, Anya Tafliovich, Jan Vahrenhold, J. Ángel Velázquez Iturbide, Eugene Wallingford, David Weintrop, Steven A. Wolfman

A formal semantics is often intended as a tool to comprehend the behavior of a language or other system. Semanticists assume, for instance, that programmers can use a semantics to understand how a particular program will behave without being forced to resort to deconstructing the output from a black-box evaluator. Indeed, different semantic models vary in what aspects of program behavior they highlight and suppress.

Every semantics has an intended audience. Formal semantics typically assume a readership with high computing or mathematical sophistication. These therefore make them inappropriate for students new to computing. What forms of description of behavior would be useful to them? In computing education, the term *notional machine* is often used to refer to a behavior description that is accessibble to beginners.

Our meeting therefore focused on what we know, and what we need to learn, about notional machines. In particular, we studied and discussed:

- Different formulations of notional machines for a variety of languages.
- The distinction between a general description of behavior, independent of a specific program, and the explication of behavior of a specific program. We argued for the value of having both the general and the specific, since learners might need to shift between the two.
- The different forms that a notional machine can take, and their styles: [MARK fill in]
- The many analogies employed in notional machines, with their respective strengths and weaknesses.
- The different forms of theories that apply to generating and understanding notional machines, including cognitive and social.
- Analogies to notional machines in other domains, from models in physics to rulebooks in board games.

We accomplished most of our stated goals: to bring together the semantics and education communities (though with much greater representation from the latter than the former); to create tutorials to educate each on the knowledge and methods of the other; and to formulate interesting examples. While there did not appear to be many long-standing "open questions", and there was not enough time to engage in editing Wikipedia, groups did organized community-wide activities (such as surveys to be conducted at upcoming conferences) and large banks of research questions (which are concrete and valuable outcomes that we had not anticipated). In sum, we believe the seminar successfully accomplished its overall stated goals.

6.43 Data Series Management

Organizers: Anthony Bagnall, Richard L. Cole, Themis Palpanas, and Konstantinos Zoumpatianos

Seminar No. 19282

Date: July 7–12, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.7.24

Participants: Azza Abouzied, Anthony Bagnall, Anastasia Bezerianos, Paul Boniol, Richard L. Cole, Michele Dallachiesa, Karima Echihabi, Jean-Daniel Fekete, Germain Forestier, Pierre Gaillard, Anna Gogolou, Søren Kejser Jensen, Mourad Khayati, Alessandro Longo, Ammar Mechouche, Abdullah Mueen, Rodica Neamtu, Themis Palpanas, John Paparrizos, Patrick Schäfer, Dennis Shasha, Nesime Tatbul, Peng Wang, Richard Wesley, Konstantinos Zoumpatianos

We now witness a very strong interest by users across different domains on data series[52] (a.k.a. time series) management systems. It is not unusual for industrial applications that produce data series to involve numbers of sequences (or subsequences) in the order of billions. As a result, analysts are unable to handle the vast amounts of data series that they have to filter and process. Consider for instance that in the health industry, for several of their analysis tasks, neuroscientists are reducing each of their 3,000 point long sequences to just the global average, because they cannot handle the size of the full sequences. Moreover, in the quest towards personalized medicine, scientists are expected to collect around 2-40 ExaBytes of DNA sequence data by 2025. In engineering, there is an abundance of sequential data. Consider for example that each engine of a Boeing Jet generates 10 TeraBytes of data every 30 minutes, while domains such as energy (i.e., wind turbine monitoring, etc.), data center, and network monitoring continuously produce measurements, forcing organizations to develop their custom solutions (i.e., Facebook Gorilla).

The goal of this seminar was to enable researchers and practitioners to exchange ideas in the topic of data series management, towards the definition of the principles necessary for the design of a big sequence management system, and the corresponding open research directions.

The seminar focused on the following key topics related to data series management:

Applications in multiple domains: We examined applications and requirements originating from various fields, including astrophysics, neuroscience, engineering, and operations management. The goal was to allow scientists and practitioners to exchange ideas, foster collaborations, and develop a common terminology.

Data series storage and access patterns: We described some of the existing (academic and commercial) systems for managing data series, examined their differences, and commented on their evolution over time. We identified their shortcomings, debated on the best ways to lay out data series on disk and in memory in order to optimize data series queries, and examined how to integrate domain specific summarizations/indexes and compression schemes in existing systems.

Query optimization: One of the most important open problems in data series management is that of query optimization. However, there has been no work on estimating the hardness/selectivity of data series similarity search queries. This is of paramount importance for effective access path selection. During the seminar we discussed the current work in the topic, and identified promising future research directions.

Machine learning and data mining for data series: Recent developments in deep neural network architectures have also caused an intense interest in examining the interactions between machine learning algorithms and data series management. We discussed machine learning from two perspectives. First, how machine learning techniques can be applied for data series analysis tasks, as well as for tuning data series management systems. Second, we how data series management systems can contribute towards the scalability of machine learning pipelines.

Visualization for data series exploration: There are several research problems in the intersection of visualization and data series management. Existing data series visualization and human interaction techniques only consider very small datasets, yet, they can play a significant role in the tasks of similarity search, analysis, and exploration of very large data series collections. We discussed open research problems along these directions, related to both the frontend and the backend.

[52] A data series, or data sequence, is an ordered set of data points.

6.44 Values in Computing

Organizers: Christoph Becker, Gregor Engels, Andrew Feenberg, Maria Angela Ferrario, and Geraldine Fitzpatrick

Seminar No. 19291

Date: July 14–19, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.7.40

Participants: Doris Allhutter, Maria Bakardjieva, Christoph Becker, Stefanie Betz, Marta Cecchinato, Teresa Cerratto-Pargman, Clarisse Sieckenius de Souza, Steve Easterbrook, Gregor Engels, Andrew Feenberg, Maria Angela Ferrario, Geraldine Fitzpatrick, Christopher Frauenberger, David Hendry, Klementina Josifovska, Selma Lamprecht, Leah Lievrouw, Ann Light, Klaus Mainzer, Curtis McCord, Lisa P. Nathan, Daniel Pargman, Elizabeth Patitsas, Austen W. Rainer, Peter Reichl, Barbara Russo, Irina Shklovski, Juliana Soares Jansen Ferreira, Sarah Spiekermann-Hoff, Dawn Walker, Blay R. Whitby, Jon Whittle, Emily Winter

The purpose of Dagstuhl Seminar 1929 'Values in Computing' was to bring together practitioners and researchers with expertise stretching beyond computer science, to include sociology, ethics, and philosophy to examine the complex relations between human values, computing technologies and society. In so doing, the seminar invited an inter-disciplinary community to share their challenges, illustrate their approaches through concrete case studies, and distil lessons learned into actionable guidelines for research and education with tangible implications for policies and industry.

The seminar was motivated by the growing urgency for computing research and industry to answer questions about the role that digital technology plays in society. The greater the scale and reach of digital technology systems, the greater their impact, both intended and unintended. Mainstream media, popular science, and the general public have only started grappling with the scale of these consequences. Many are calling institutions, professionals, and scientists to act [3]. Recent years have seen an increasing number of high-profile software scandals and malpractices in which individual privacy and democracy have been undermined (Cambridge Analytica's use of Facebook data), the environmental impact of air pollutants disregarded (the Volkswagen's diesel emission scandal), and human lives lost (the Boeing 737 Max anti-stall software disasters).

These events are the constant reminders that human values are indeed "the facts of the future" [1], as Feenberg argues. Values are not the opposite of facts, they become facts: the more weight we give to certain values (e.g. wealth, political influence, power), the bigger the 'blind spots' of the existing values become (e.g. environmental sustainability, equality and social justice). There is a pressing need then to understand how human values operate and to build on this understanding to consider how research and education might contribute to a more socially responsible computing industry.

To this end, the seminar brought together disciplines with a long tradition of critical thinking and human-centred approaches to computing with those that, such as Software Engineering, have been traditionally considered, albeit increasingly controversially, as 'values neutral'. The breadth and depth of the interdisciplinary debate, one of the key distinguishing features and strengths of this five-day seminar, was also, and intentionally so, one of its main challenges. This was particularly evident when the need to unpack the multifaceted and often abstract notion of human values was met by the demand for the discussion to be of concrete relevance to computing education and practice. Within this context, one of the key objectives of the seminar was to facilitate both the exploration of broad themes and the identification of specific topics that would require meaningful cross-disciplinary effort. To this end, a two-pronged approach was designed to encourage both divergence and converge of viewpoints.

Thematic divergence was encouraged through six short *Seed Talks*, ten open-floor *Lightning Talks*, and a *Soap Box* session were participants would pitch high-level challenges to provoke discussion. Convergence was facilitated by *World Café* style group discussions around six emergent themes. Over the last two days, these themes were then distilled into four topics with one working group assigned to each (*Action*, *Education*, *Research*, and *Response*). Seed Talks were invited 20-minute talks designed to be informative and provocative. Thematically, they were structured around the original seminar proposal scoping areas: theory and methodologies (Feenberg and Mainzer), professional practice (Spiekermann and Whittle); and educational pathways (Nathan and Patitsas). Participants offered Lightning Talks on a variety of topics of their own choosing. For instance, Easterbrook focused on the environmental crisis and called for urgent action; Walker, McCord and Lievrouw shared their experiences of socially responsible digital activism; Frauenberger provided concrete examples on how different ways of thinking informatics

in education [2]; Winter outlined the tools and techniques used to study values in software production [4]; and Jensen-Ferreira shared her approach to software industry research. Finally, four teams worked on a specific *Values in Computing* topic, each identifying a possible course of action:

1. *Action* – This group worked under the premise that the professional knowledge and critical insight of computer and social scientists should be mobilized as an active force in public education and policy-making concerning the design, implementation and regulation of information technology. With a view to these three lines of action, the group proposed the penning and wide distribution of a document, tentatively entitled *"The Dagstuhl Declaration"* here included.
2. *Research* – The Research group pursued a threefold-goal: understand the state-of-the-art of the research and highlight under-explored research areas; discuss methods and tools that have been or can be used, and identify future research directions.
3. *Education* – The Education group discussed the implications for undergraduate and graduate computing education by conducting a brief but focused exploration of existing university-level courses, methods and tools and their mapping of curriculum cross-cutting learning objectives.
4. *Response* – This group worked on the intersection between climate emergency and the future of computing and centred its activity on gathering resources about this intersection and writing an opinion piece to address it.

References

1 Andrew Feenberg. Ten paradoxes of technology. Techné: Research in Philosophy and Technology, 14(1):3–15, 2010.

2 Christopher Frauenberger and Peter Purgathofer. 2019. Ways of thinking in informatics. Commun. ACM 62, 7 (June 2019), 58-64.

3 Leon J. Osterweil. Be prepared. SIGSOFT Softw. Eng. Notes, 41(5):4–5, November 2016.

4 Emily Winter, Stephen Forshaw, Lucy Hunt, and Maria Angela Ferrario. 2019. Towards a systematic study of values in SE: tools for industry and education. In Proceedings of the 41st International Conference on Software Engineering: ICSE-NIER '19. IEEE Press, Piscataway, NJ, USA, 61-64.

6

6.45 Mobile Data Visualization

Organizers: Eun Kyoung Choe, Raimund Dachselt, Petra Isenberg, and Bongshin Lee

Seminar No. 19292

Date: July 14–19, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.7.78

Participants: Wolfgang Aigner, Dominikus Baur, Frank Bentley, Lonni Besancon, Tanja Blascheck, Matthew Brehmer, Sheelagh Carpendale, Eun Kyoung Choe, Christopher Collins, Raimund Dachselt, Tim Dwyer, Daniel Epstein, Tom Horak, Pourang P. Irani, Petra Isenberg, Tobias Isenberg, Alark Joshi, Ricardo Langner, Bongshin Lee, Lena Mamykina, Charles Perin, Harald Reiterer, John T. Stasko, Christian Tominski, Jo Vermeulen

As pen- and/or touch-enabled mobile devices have become more powerful and ubiquitous, we see a growing demand for *mobile data visualization* to facilitate visual access to data on mobile devices (see Figure 6.10 for examples). Lay people increasingly access a wide range of data, including weather, finance, and personal health on their phone. Small business owners start to use business intelligence software equipped with data visualization on mobile devices to make better business decisions. In responding to these needs, practitioners have actively been designing mobile visualizations embedded in commercial systems. However, research communities, such as Human-Computer Interaction (HCI), Information Visualization (InfoVis), and Ubiquitous Computing (UbiComp) have not paid enough attention to mobile data visualization.

Over the past few decades, the visualization research community has conducted extensive research, designing and developing a large number of visualization techniques and systems mostly for a desktop environment. However, the accumulated knowledge may not be readily transferable to mobile devices due to their fundamental differences in their display size, interaction, and target audience, among others. The small display on mobile devices is more vulnerable to the scalability issue and poses a well-known challenge, the fat finger problem. Mouse-over interaction, which is prevalent in interactive visualization systems in the desktop environment, is not available on mobile devices. While traditional visualizations mainly target data-savvy groups of people such as scientists and researchers, visualizations on mobile devices should account for a broader range of target audience, including lay people who might have low data and visualization literacy.

Fig. 6.10
Examples of mobile data visualizations: step count and sleep data visualization on Fitbit Ionic and mobile app (top left); a multiple coordinated views across two mobile devices in VisTiles (bottom left); and visual data exploration on a tablet leveraging pen and touch interaction in TouchPivot (right).

This Dagstuhl seminar follows in the footsteps of the "Data Visualization on Mobile Devices" workshop at CHI 2018, our initial effort in establishing a community around mobile data visualization. We brought researchers and practitioners from relevant application and research fields, including InfoVis, UbiComp, mobile HCI, and interaction design to exchange information and experiences, to stimulate discussion, to make new connections, and to identify novel ideas and future directions around mobile data visualization.

Unlike the CHI workshop, this five-day Dagstuhl workshop enabled us to explore mobile data visualization in depth through speedy & intense research exchanges, interactive demos & tutorials, as well as active breakout group discussions.

The Week at a Glance

Monday. The seminar was kicked off by the organizers with an introduction to the topic of mobile data visualization and by providing organizational information. Afterwards, all participants introduced themselves and their expectations with a short two-minute slide presentation. This session was followed by a speedy research brainstorming activity (see Figure 6.11): In rapid five-minute sessions, two participants facing each other introduced their research activities and jointly sketched new ideas. By rotating half of the group, each session was repeated eleven times with new constellations of two people each time.

Fig. 6.11
Exchange of research interest & background in a speed dating format.

In the afternoon, five demo stations were set up and participants were split into groups to attend them in turn. Five researchers presented their latest mobile visualization demos in hands-on sessions (see Figure 6.12). These were:

- Tanja Blascheck: Smartwatch demo from a study comparing three representations–bar, donut, text (joint work with Lonni Besançon, Anastasia Bezerianos, Bongshin Lee, Petra Isenberg).
- Matthew Bremer: Tilting, brushing, & dialing for mobile vis (joint work with Bongshin Lee, Christopher Collins, Ken Hinckley).
- Tobias Isenberg: Personal home automation system with mobile data access and control.
- Alark Joshi: Visualization of off-screen data using summarization techniques (joint work with Martino Kuan, Alejandro Garcia, Sophie Engle).
- Jo Vermeulen: Product Fingerprints, a mobile visualization that allows people to compare nutritional information between food products (joint work with Carrie Mah, Kevin Ta, Samuel Huron, Richard Pusch, Jo Vermeulen, Lora Oehlberg, Sheelagh Carpendale).

Fig. 6.12
One of the mobile visualization demos presented to a small group of participants.

In a second activity, 14 participants presented a design critique of an existing mobile visualization, partly commercial products, partly research results (see Figure 6.13). Besides evoking the spirit of a good discussion, it helped getting a broad overview about currently available solutions.

Fig. 6.13
Impressions from the Design Critique Session.

In a followup activity, to arrive at a common understanding of the state of the art in mobile data visualization, we split attendees into three groups according to their main expertise. The three groups were:

- Information Visualization–Mobile Visualization Resources
- Visualization in Ubiquitous Computing Research
- Mobile Interaction and Human Computer Interaction

Each group was tasked to collect and discuss the state of the art, with an end goal of creating a short presentation to be given to the entire audience. As a result, the collected material and insights were presented to the plenum by each group.

Through these diverse activities during the first day, participants did not only gain a good understanding of each other's background and research interests, but also established a common ground and expertise in the field of mobile data visualization

Tuesday. The second day started with a lively brainstorming and discussion of challenges and important research questions in the field of mobile data visualization. From about ten larger topics we identified, four were chosen to form parallel breakout groups:

- Group 1: Evaluating Mobile Data Visualization
- Group 2: What is Mobile Vis?
- Group 3: Responsive Visualization
- Group 4: Vis for Good & Ethics

Using the impressive facilities of Dagstuhl in terms of rooms and places, space to think and coffee to drink, we had intense discussions within each group. We generated deeper research questions and challenges, and identified collaborative cross-disciplinary research opportunities and approaches. Section 4 provides more details on each of these working groups.

After lunch, groups reported back on what they had discussed (see Figure 6.14). The four groups decided to continue and deepen their discussions in the afternoon, this time focusing more on what could become a concrete research output.

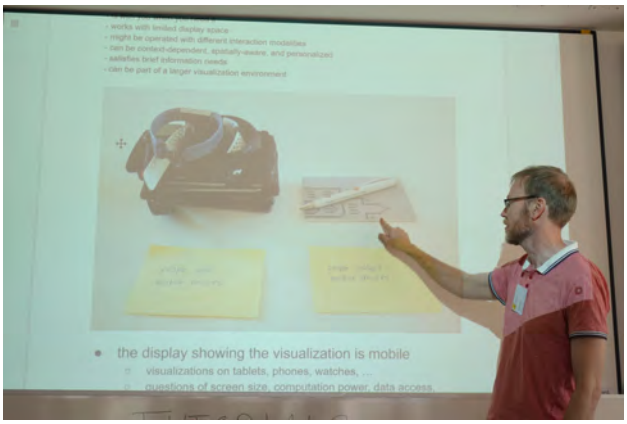

Fig. 6.14
Report back from Group 2 on "What is Mobile Vis?"

Wednesday. Wednesday morning was devoted to the presentation of tutorials. Five participants had volunteered to give tutorials in two time slots, allowing other participants to attend two one-hour tutorials. Figure 6 shows the title slides of all informative and well-received tutorials, and Section 3 provides details on each of them.

Following the tradition of Dagstuhl Seminar, Wednesday afternoon was set aside for social activities. We took the bus to experience the famous Saarschleife high from the impressive treetop walk. Visiting Mettlach and having dinner in a brewery intensified personal conversations and fostered planning for joint research collaborations.

Thursday. Similar to Tuesday, the entire day was dedicated to group work (see Figure 7). The list of possible topics for breakout groups was revisited, and people assembled to form new groups on other challenging topics:

- Group 5: Starting Mobile Visualization from Scratch
- Group 6: Beyond Watch & Phone: From Mobile to Ubiquitous Visualization
- Group 7: (Discoverable) Interaction for Mobile Visualization
- Group 8: From Perception to Behavior Change: Designing and Evaluating Glanceable Mobile Vis
- Group 9: Mobile Vis for 3D Data / AR Vis

Again, both the morning and afternoon were used for intensely discussing challenges, defining design spaces, shaping the knowledge on the given topic, and identifying opportunities for joint research. Groups also reported back to the plenum, and results were discussed openly. Section 4 provides more details on each of these working groups and their outcomes.

Friday. After interesting and enriching days of joint discussions, which considerably broadened the horizon, time had come to wrap up the seminar on Mobile Data Visualization. Most importantly, a broad range of future collaborative activities were discussed: writing a state-of-the-art report, joint grant proposals, further workshop proposals, individual papers, editing a special journal issue, and writing a book on the topic. In the end, we agreed on a book as a possible major outcome (see Section 5). Organizational details were clarified, before the seminar was concluded with thanking all participants for their great contributions and commitment during the entire week.

6.46 Secure Composition for Hardware Systems

Organizers: Divya Arora, Ilia Polian, Francesco Regazzoni, and Patrick Schaumont

Seminar No. 19301

Date: July 21–26, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.7.94

Participants: Divya Arora, Georg T. Becker, Yaacov Belenky, Shivam Bhasin, Ran Canetti, Gaetan Cassiers, Anupam Chattopadhyay, Jean-Luc Danger, Lucas Davi, Elke De Mulder, Elena Dubrova, Yunsi Fei, Fatemeh Ganji, Tim Erhan Güneysu, Annelie Heuser, Johann Heyszl, Elif Bilge Kavun, Osnat Keren, Johann Knechtel, Itamar Levi, Michail Maniatakos, Marcel Medwed, Nele Mentens, Johannes Mittmann, Debdeep Mukhopadhyay, Paolo Palmieri, Ilia Polian, Milos Prvulovic, Wenjing Rao, Francesco Regazzoni, Ahmad-Reza Sadeghi, Kazuo Sakiyama, Fareena Saqib, Patrick Schaumont, Werner Schindler, Georg Sigl, Dey Soumyajit, Marc Stöttinger, Shahin Tajik, Marten Van Dijk, Ingrid Verbauwhede

Today's electronic systems consist of mixtures of programmable, reconfigurable and application-specific hardware components, tied together by tremendously complex software. At the same time, systems are increasingly integrated such that a system that was traditionally regarded "harmless" (e.g., an entertainment system in a car) finds itself tightly coupled with safety-critical driving-assistance systems and security-sensitive online payment systems. Moreover, a system's hardware components are now often directly accessible to the users, making the system vulnerable to physical attacks via its hardware which becomes the system's "Achille's heel". This necessitates a new look on system security from hardware perspective.

The Dagstuhl seminar "Secure Composition for Hardware Systems", which took place on July 21-26, 2019, focused on secure composition of systems which contain hardware blocks. This is a practically important but a theoretically challenging problem where several foundational questions still lack an adequate answer.

Several formats were used during the seminar. The first phase of the seminar, which focused on prior findings, started with presentations by five pre-selected experts giving their view on secure composition from different perspectives: theory, design automation, trusted execution environments and attacks countermeasures. Then, small-group discussions of relevant state of the art were held, focusing on questions such as "What does it mean to securely compose two elements?" or "What is the role of models in secure composition?" The findings of the small groups were intensively discussed in plenary sessions.

The second phase of the seminar was devoted to discussing research questions. Some of the questions were prepared by the seminar organizers (e.g., "Which protocol-level secure composition methods are applicable in hardware domain?" or "How to counter possible loss of security due to abstraction of hardware components?") and some additional questions were proposed by the participants (e.g., "How to bootstrap trust in a distributed hardware system?"). The questions were discussed again in small groups, intertwined by individual presentations in plenum (for instance, an in-depth study on the applicability of Universal Composability (UC) in the hardware domain).

Two immediate outcomes grew out of the seminar. First, some participants are organizing a special session on secure compositions in one of the leading scientific conferences; a respective proposal was recently accepted by the "Design, Automation, and Test in Europe Conference" (DATE). Second, there is an ambitious plan to prepare a manuscript on the full variety of aspects in secure composition of electronic systems and submit it as a "Systematization of Knowledge" (SoK) paper to the IEEE Symposium on Security and Privacy (S&P); this effort is ongoing at the time of writing this report.

Overall, we believe that this seminar has provided entirely new insights to most of the participants and has opened new avenues for research on the intersection of security and hardware systems. It brought together researchers from communities who rarely interacted with each other in the past. The seminar helped define new research challenges, and activities are underway to put the topic of secure composition higher on the agenda of the respective communities.

The organizers are thankful to the Dagstuhl team (and in particular to Dr. Andreas Dolzmann who handled the scientific part and Mrs. Heike Clemens who was of invaluable help in organizing the social event and masterly handled all practical issues); to Dr. Elif Bilge Kavun who did a great job in collecting and organizing the documents from participants and in preparing the summarizing texts; and to all the participants for making this seminar a success.

6

6.47 Cybersafety Threats – from Deception to Aggression

Organizers: Zinaida Benenson, Marianne Junger, Daniela Oliveira, and Gianluca Stringhini

Seminar No. 19302

Date: July 21–26, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.7.117

Participants: Abhishta Abhishta, Zinaida Benenson, Matt Bishop, Jan-Willem Bullée, Joe Calandrino, Deanna Caputo, Claude Castelluccia, Yi Ting Chua, Natalie Ebner, Matthew Edwards, Manuel Egele, Jeremy J. Epstein, Freya Gassmann, Alice Hutchings, Marianne Junger, Katsiaryna Labunets, Elmer Lastdrager, Gabriele Lenzini, Daniela Oliveira, Simon Parkin, William Robertson, Stefan Schiffner, Michael Sirivianos, Ivan Srba, Gianluca Stringhini, Gareth Tyson, Sophie van Der Zee, Sebastian Wachs, Victoria Wang, Jeff Yan, Savvas Zannettou

A number of malicious activities are prospering online and are putting users at risk. In particular, cyber deception and cyber aggression practices are increasing their reach and seriousness, leading to a number of harmful practices such as phishing, disinformation, radicalization, and cyberbullying. Attack strategies include controlling and operating fake or compromised social media accounts, artificially manipulating the reputation of online entities, spreading false information, and manipulating users via psychological principles of influence into performing behaviors that are counter to their best interests and benefit the attackers.

So far, computer science research on cybersafety has looked at the various sub-problems in isolation, mostly relying on algorithms aimed at threat detection, and without considering the implications of the attacks and countermeasures for individual users as well as for society. On the other hand, human factors and social science researchers often consider user interfaces and social interactions without taking full advantage of the algorithmic, data-driven cybersafety research. Moreover, the legal and ethical implications of attacks and countermeasures are often unclear.

The goal of the Dagstuhl Seminar 19302 "Cybersafety Threats – from Deception to Aggression" was to provide a platform for researchers to look at the problem of cybersafety from a holistic and multi-disciplinary perspective. The participants were drawn from a number of disciplines such as computer science, criminology, psychology, and education, with the aim of developing new ideas to understand and mitigate the problems.

At the beginning of the seminar, we asked participants to identify important themes to focus on, and these themes were refined through specific activities and discussions during the first day: Firstly, all participants gave 5-minute talks where they presented their current research related to the seminar, and their expectations and topics they would like to work on during the week. Secondly, we conducted three *introductory panels* on the topics of *Cyber Deception*, *Cyber Aggression* and *Propaganda & Disinformation*. Each panel consisted of five participants. We took special care to represent different disciplines and different career stages in each panel.

By the beginning of the second day, participants had identified four key themes to study in this area, which we describe in detail in the rest of this section. The participants formed working groups (WGs) for each theme.

Theme 1: Attacker modeling

The working group focused on predicting the next steps of an ongoing attack by means of a probabilistic model. The initial model developed by the group consists of 9 variables: attacker goals, characteristics of the attack (e.g., how long the attack takes, tools employed), consequences, authorization, attribution, expected resilience of the victim, expected characteristics of the victim from attacker's perspective, actual characteristics of the victim, actual responsiveness of the victim. The developed model was verified and refined using two known attacks as case studies: the Internet Worm (1988) and the SpamHaus DDoS attack (2013).

Two most important next steps to refine the model are:

1. Convert the variables into measurable quantities
2. Obtain labeled data on which the model can be trained

The working group started working on a conceptual paper that describes the model, and discussed possible venues for its publication. Several methods of obtaining the data for the model were proposed, such as interviewing CISOs and other defenders, creating financial incentives for organization to share their data, and organizing a stakeholder workshop including not only defenders, but also former attackers who now work as security consultants.

Theme 2: Unintended consequences of countermeasures

This working group focused on an often overlooked aspect of computer security research: the fact that deploying any countermeasure to mitigate malicious online activity can have unexpected consequences and harms to other parties. The members of this working group started by discussing a number of scenarios: intimate partner abuse, CEO fraud, disinformation, online dating fraud, and phishing, and developed a taxonomy of these potential harms. The taxonomy takes into account not only technical issues that might arise from deploying countermeasures but also socio-technical ones such as the displacement effect of attackers moving to other victims, the additional costs incurred by using the countermeasure, and the issues arising from complacency, for example leaving users desensitized by displaying too many alerts to prevent a certain type of attack.

Theme 3: Measuring human behavior from information security (and societal) perspectives

Measuring online behavior is of fundamental importance to gain an accurate understanding of malicious online activities such as cybercrime. The research community, however, does not have well established techniques to accurately measure this behavior, and this can lead to studies presenting largely contradicting results. This working group focused on identifying techniques relevant to measure and model various types of online behavior, from cyberbullying and disinformation to ransomware and phishing. As a final outcome, the working group drafted two methodological frameworks for researchers aiming to study these problems, one focused on socio-technical threats (cyberbullying and disinformation) and one focused on cybersecurity (phishing and malware).

Theme 4: Prevention, detection, response and recovery.

A key challenge when mitigating socio-technical issues is developing the most effective countermeasures. This group focused on developing detection and prevention approaches focusing on threats encountered by adolescents when surfing the Web (e.g., cybergrooming). A common issue here is that adolescents rarely turn to adults for help, and therefore any mitigation based on direct parental oversight has limited effectiveness. To go beyond these issues, the group developed a mitigation strategy based on a "guardian angel" approach. The idea is to let a minor create a "guardian avatar" that will then advise them on cybersafety practices, with a decreasing level of oversight as the minor grows up. While the children are very young, the guardian avatar will closely supervise them, reporting any suspicious contacts that they have online to a parent or a guardian. Later, as the child enters adolescence, the avatar will gradually take on an advisory role, eventually only providing advice once the adolescent asks for it. The group considered privacy issues and interdisciplinary aspects related to psychology and education, and developed a proposal of how the avatar would work.

Conclusion and Future Work

The seminar produced a number of ideas on how to investigate and mitigate cybersafety threats. It enabled researchers from different disciplines to connect, and set the agenda for potentially impactful research to be carried out in the next years. Joint publications and funding for joint research were discussed in each WG and later in the plenum. For example, WG 3 considered possibilities for a large international grant, such as H2020. The ideas produced as part of theme 4 resulted in the paper "Identifying Unintended Harms of Cybersecurity Countermeasures" to appear at the APWG eCrime Symposium in November 2019.

6.48 Software Protection Decision Support and Evaluation Methodologies

Organizers: Bjorn De Sutter, Christian Collberg, Mila Dalla Preda, and Brecht Wyseur

Seminar No. 19331

Date: August 11–16, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.8.1

Participants: Mohsen Ahmadvand, Sébastien Bardin, Cataldo Basile, Tim Blazytko, Richard Bonichon, Richard Clayton, Christian Collberg, Moritz Contag, Bart Coppens, Jorge R. Cuéllar, Mila Dalla Preda, Bjorn De Sutter, Laurent Dore, Ninon Eyrolles, Roberto Giacobazzi, Yuan Xiang Gu, Christophe Hauser, Stefan Katzenbeisser, Eric Lafortune, Peter Lafosse, Patrik Marcacci, J. Todd McDonald, Christian Mönch, Leon Moonen, Jan Newger, Katharina Pfeffer, Yannik Potdevin, Uwe Resas, Rolf Rolles, Sebastian Schrittwieser, Bahman Sistany, Natalia Stakhanova, Atis Straujums, Stijn Volckaert, John Wagner, Andreas Weber, Brecht Wyseur, Michael Zunke

Overview and Motivation

The area of Man-At-The-End (MATE) software protection is an evolving battlefield on which attackers execute white-box attacks: They control the devices and environments and use a range of tools to inspect, analyze, and alter software and its assets. Their tools include disassemblers, code browsers, debuggers, emulators, instrumentation tools, fuzzers, symbolic execution engines, customized OS features, pattern matchers, etc.

To meet the security requirements of assets embedded in software, i.e., valuable data and code, many protections need to be composed. Those requirements include the confidentially of secret keys and software IP (novel algorithms, novel deep learning models, ...), and the integrity of license checking code and anti-copy protections. Attackers attack them through reverse engineering and tampering, for which they use the aforementioned tools and for which they often can afford spending time and effort on executing many, highly complex and time-consuming, manual and automated analyses. The need for composing many protections follows from the fact that advanced attackers can use all the mentioned tools and try many different approaches. In other words, to be effective, the deployed protections need to protect against all possible attack vectors.

As all protections come with overhead, and as many of them have downsides that complicate various aspects of the software development life cycle (SDLC), the users of a software protection tool cannot simply deploy all available protections. Instead, they have to select the protections and their parameters for every single asset in a program, taking into account non-functional requirements for the whole program and its SDLC.

The organizers of this workshop, and many experts in their network, consider the lack of automated decision support for selecting the best protections, and the lack of a generally accepted, broadly applicable methodology to evaluate and quantify the strength of a selected combination, the biggest challenges in the domain of software protection. As a result, the deployment of software protection is most often not trustworthy, error-prone, not measurable, and extremely expensive because experts are needed and they need a lot of time, increasing the time to market.

This situation is becoming ever more problematic. For example, connected intelligent vehicles are quickly being deployed in the market now and autonomous vehicles are going to be deployed in 3-5 years. Software protection evaluation and measurement research and development must match up that pace to provide enough technology support for controllable and scientific methods to manage the quality of automotive security as key part of vehicle reliability and safety. There is hence a huge need to make progress w.r.t. software protection decision support and evaluation methodologies, the topic of the proposed seminar.

Goals of the Seminar

Following a pre-seminar survey among the registered participants to focus the seminar and to select the highest priority objectives among the many possible ones, the primary goal of the seminar was determined to be the foundations of a white paper on software protection evaluation methodologies, to be used as a best practices guideline by researchers and practitioners when they evaluate (combinations of) defensive and/or offensive techniques in the domain of MATE software protection. This can also serve as a guideline to reviewers of submitted journal and conference papers in which novel techniques are proposed and evaluated. A secondary goal was the establishment of good benchmarking practices, including the choice of suitable benchmarks and the selection and generation thereof for use in future research in MATE software protection. A third goal was to collect feedback and ideas on how to push the state of the art in decision support systems.

Week Overview

Preparation. Prior to the seminar, the organizers set up a survey to collect the necessary information for a seminar bundle that provided background information about and to all participants. Moreover, they collected information regarding the potential outcomes that participants were most interested in, to which ones they could likely contribute, and which potential outcomes they considered most likely to make progress on. Furthermore, a reading list was presented to the participants with the goal of getting everyone on the same page as much and as soon as possible [1–8].

Whereas the schedule for the first two days was mostly fixed a priori, the schedule for later days was more dynamic, as it was adapted to the feedback obtained by the organizers during the early days, and to the outcomes of different sessions.

Monday. The first day was devoted to setting the scope of the seminar, and clarifying the seminar goals, strategy, and plan. In the morning, three overviews were presented of man-at-the-end software protection techniques in the scope of the seminar, as well as some attacks on them. These presentations focused on obfuscation vs. static analysis, (anti-)tampering in online games, and additional protections beyond the ones discussed in the first two presentations.

In the early afternoon, four deeper technical introductions were presented of four more concrete classes of defensive and corresponding offensive techniques that would serve as case studies throughout the seminar: 1) virtual machine obfuscation, 2) (anti-)disassembly, 3) trace semantics based attacks, and 4) data obfuscation. The strategy for the week was to brainstorm about these concrete techniques first, in particular on how the strength of these techniques are supposed to be evaluated, e.g., in papers that present novel (combinations of) techniques, or in penetration tests. Later, the concrete results for the individual case studies would then be generalized into best practices and guidelines for software protection evaluation methodologies.

Whereas the morning presentations and most of the case studies focused mostly on defensive techniques, three presentations in the afternoon provided complementary insights about offensive techniques, ranging from more academic semantics-based attack techniques, over an industrial case study of deobfuscation of compile-time obfuscation, and offensive techniques in binary analysis.

Thus, the scene was set in terms of both defensive and offensive techniques, and all participants to a large degree spoke the same language before starting the brainstorm sessions in the rest of the week.

Tuesday. Tuesday focused mostly on the seminar track of software protection evaluation methodologies.

In the early morning, additional input was provided on existing, already studied aspects relevant to such methodologies. This included software protection metrics, empirical experiments to assess protections, and security economics. These presentations provided useful hooks for the next session, which consisted of parallel, small break-out brainstorm sessions (three groups per case study) on the first two case studies. In these brainstorm sessions, the goal was to provide answers to questions such as the following:

- What would a document similar to the SIGPLAN empirical evaluation checklist look like for papers presenting new VM-based protections?
- Which requirements or recommendations can we put forward with respect to the protected objects (i.e., benchmarks) and their treatment (i.e., how they are created, compiled, ...) for the evaluation?
- What aspects of the attack models and which assumptions should be made explicit, which ones should be justified, e.g., regarding attacker goals and attacker activities.
- How should sensitivity to different inputs (e.g., random generator seeds, configuration options, features of code samples, ...) be evaluated and discussed?
- What threats to validity should be discussed?
- What aspects of the protection should be evaluated (potency, resilience, learnability, usability, stealth, renewability, different forms of costs, ...)?
- Under what conditions would you consider the protection to be "real world" applicable?
- What flaws (e.g., unrealistic assumptions) have you seen in existing papers that should be avoided?
- What are (minimal) requirements / recommendations regarding reproducibility?
- What pitfalls can you list that we should share with people?

After the independent brainstorms in small groups and following lunch, the three groups per case study came together to merge the results of their brainstorms, after which the merged results were shared in a plenary session.

Later in the afternoon, additional ideas were presented on topics relevant for software protection evaluation methodologies. The covered topics were benchmark generation, security activities in protected software product life cycles, the resilience of software integrity protection (work in progress), and a (unified) measure theory for potency. These topics were presented after the initial brainstorms not to bias those brainstorms. Their nature was more forward looking, covering a number of open challenges as well as potential directions for future research. They offered the speakers a sound board to get feedback and could serve as the starting point of informal discussions later in the seminar.

While the practice is discouraged by the Dagstuhl administration, we still decided to organize an evening session on Tuesday. Afterwards, we realized that this made the seminar a bit too dense, but it did serve the useful purpose of introducing the participants to the seminar track on decision support tools for software protection early enough in the seminar to allow enough time for informal discussions with and between researchers active on this topic during the remainder of the week. This was especially useful to allow those academic researchers to check the validity of some of their assumptions about real-world aspects with the present practitioners from industry and with researchers from other domains.

Besides an overview of an existing design and implementation of a software protection decision support system, a hands-on walk through of a practical attack on a virtual machine protection (as in one of the case studies) was presented, as well as some ideas to make such protection stronger.

[53] For some reason, most of us don't remember the rest of the evening in enough detail to report on it reliably.

Wednesday. Early on Wednesday morning, the focus shifted towards decision support tools, with three presentations by practitioners in companies that provide software protection solutions. These presentations focused on the support they provide to help their customers use their tools.

Later in the morning, case studies 3 and 4 were discussed in another round of parallel, small group break-out brainstorm sessions.

In the afternoon, the social outing took place, which consisted of a visit to Trier and a wine tasting at a winery where we also had dinner.[53]

Thursday. On Thursday morning, another round of break-out sessions was organized to structure the outcomes of the first round. Based on inputs collected during the first three days, the organizers drafted a structure for a white paper on software protection methodologies. In 4 parallel sessions, the participants brainstormed on how to fit the results of the first round (i.e., bullet points with concrete guidelines and considerations for each case study) into that structure, and which parts of those results could be generalized beyond the individual case studies. In a plenary session, the results of these break-outs were then presented.

In addition, the specific topic of benchmarking was discussed, focusing on questions regarding the required features of benchmarks (e.g., should or should they not contain actual security-sensitive assets) as well as potential strategies to get from the situation today, in which very few benchmarks used in papers are available for reproducing the results, to a situation in which a standard set of benchmarks is available and effectively used in studies.

In the afternoon, several demonstrations of practical tools were given, including the already mentioned decision support system of which the concepts had been presented on Tuesday evening and the Binary Ninja disassembler that is rapidly gaining popularity. Two presentations were also given on usable security and challenges and capabilities of modern static analysis of obfuscated code. There provided additional insights useful for both designers of decision support tools and evaluation methodologies.

Friday. The last morning started off with a potpourri of interesting topics that did not fit well in the main tracks of general evaluation methodologies and decision support on the one hand, and benchmarking on the other. Given the availability of many experts in the domain of software protection, we decided that everyone that wanted to launch new ideas or collect feedback on them in the broad domain of the seminar should have that chance. So the day started with short presentations on the protection of machine learning as a specific new type of application, on security levels for white-box cryptography, and on hardware/software binding using DRAM.

Later in the morning, the seminar was wrapped up with a discussion of the outcomes so far, and an agreement on plans to continue the work on the software protection evaluation methodology white paper and the assembly of a benchmark collection.

References

1 S. Schrittwieser, S. Katzenbeisser, J.Kinder, G. Merzdovnik, and E. Weippl: Protecting software through obfuscation: Can it keep pace with progress in code analysis? *ACM Comput. Surv.*, **49**(1), 2016.

2 M. Ceccato, P. Tonella P, C. Basile, P. Falcarin, M. Torchiano, B. Coppens, and B. De Sutter: Understanding the behaviour of hackers while performing attack tasks in a professional setting and in a public challenge. *Empirical Software Engineering* 2018; **24**(1):240–286.

3 B. Cataldo, D. Canavese, L. Regano, P. Falcarin, and B. De Sutter: A Meta-model for Software Protections and Reverse Engineering Attacks. *Journal of Systems and Software* 150 (April): 3–21, 2019

4 B. Yadegari, B. Johannesmeyer, B. Whitely, and S. Debray: A generic approach to automatic deobfuscation of executable code. In: *Proc. IEEE Symposium on Security and Privacy*, pp. 674–691 (2015)

5 T. Blazytko, M. Contag, C. Aschermann, and T. Holz: Syntia: synthesizing the semantics of obfuscated code. *Proc. of the 26th USENIX Security Symposium (SEC'17)*, pp. 643–659. 2017

6 S. Banescu, C. Collberg, and A. Pretschner: Predicting the Resilience of Obfuscated Code Against Symbolic Execution Attacks via Machine Learning. *Proc. of the 26th USENIX Conference on Security Symposium (SEC'17)*, pp. 661-678, 2017

7 C. Basile et al.: D5.11 ASPIRE Framework Report. Technical Report ASPIRE project. https://aspire-fp7.eu/sites/default/files/D5.11-ASPIRE-Framework-Report.pdf

8 M. Ceccato et al.: D4.06 ASPIRE Security Evaluation Methodology – Security Evaluation. Technical Report ASPIRE project. https://aspire-fp7.eu/sites/default/files/D4.06-ASPIRE-Security-Evaluation-Methodology.pdf

6.49 Algorithms and Complexity for Continuous Problems

Organizers: Dmitriy Bilyk, Aicke Hinrichs, Frances Y. Kuo, and Klaus Ritter

Seminar No. 19341

Date: August 18–23, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.8.26

Participants: Dmitriy Bilyk, James M. Calvin, Ronald Cools, Sonja Cox, Steffen Dereich, Benjamin Doerr, Martin Ehler, Michael Giles, Michael Gnewuch, Takashi Goda, Mario Hefter, Stefan Heinrich, Aicke Hinrichs, Martin Hutzenthaler, Lutz Kämmerer, Alexander Keller, Kristin Kirchner, David Krieg, Peter Kritzer, Robert J. Kunsch, Frances Y. Kuo, Gunther Leobacher, Thomas Müller-Gronbach, Erich Novak, Dirk Nuyens, Friedrich Pillichshammer, Leszek Plaskota, Joscha Prochno, Pawel Przybylowicz, Klaus Ritter, Daniel Rudolf, Stefan Steinerberger, Michaela Szölgyenyi, Aretha Teckentrup, Vladimir N. Temlyakov, Mario Ullrich, Jan Vybíral, Marcin Wnuk, Henryk Wozniakowski, Larisa Yaroslavtseva, Marguerite Zani

This was already the 13th Dagstuhl Seminar on Algorithms and Complexity for Continuous Problems over a period of 28 years. It brought together researchers from different communities working on complexity of continuous problems. Such problems, which originate from numerous areas, including physics, chemistry, finance, and economics, can almost never be solved analytically, but rather only approximately to within some error threshold. The complexity analysis ideally includes the construction of (asymptotically) optimal algorithms. Although the seminar title has remained the same, many of the topics and participants change with each seminar and each seminar in this series is of a very interdisciplinary nature. The current seminar attracted 41 participants from nine different countries all over the world. About 30% of them were young researchers including PhD students. There were 34 presentations.

The following topics were covered:

Tractability analysis of high-dimensional problems: Tractability analysis is an area of applied mathematics and theoretical computer science that studies the minimal computational resources needed for the approximate solution of problems with a huge number of variables, and it can be seen as a unifying theme for the preceding seminars in this series. Many concrete problems from applications have been analyzed in this context, new algorithms were developed, approaches to break the curse of dimensionality were established, but there remain a number of important open problems. Tractability analysis will serve as a guideline and a tool for establishing complexity results and for constructing algorithms for infinite dimensional problems.

Computational stochastics: The focus was on weak and strong approximation as well as on the quadrature problem for stochastic ordinary or partial differential equations, i.e., on models with a random dynamics in a finite- or infinite-dimensional state space. A major topic was the complexity analysis for stochastic differential equations under non-standard assumptions.

Computing and complexity in infinite dimensions: Computational problems with infinitely many variables naturally arise in rather different application areas. Results and techniques from tractability analysis are available and thus permit one to study infinite dimensional problems as the limit of finite dimensional ones. Moreover, the availability of generic types of algorithms, like the multivariate decomposition method or the multi-level approach, will contribute to the complexity analysis and practical application in integration and approximation problems of infinitely many variables.

Discrepancy theory: Classical discrepancy theory is concerned with the question how uniformly finite point sets can be distributed. The geometric notion of discrepancy is intimately connected to the complexity of integration for functions from certain function classes. For problems in both fixed low dimension and high dimension, there are intriguing open questions whose solution would impact both fields of discrepancy theory and tractability studies.

Computational/applied harmonic analysis: Harmonic analysis plays an increasingly important role both in discrepancy theory and tractability analysis. One highlight is the proof of the currently best known lower bound for the star discrepancy in fixed dimension, which showed close connections between different areas, so similar techniques could be used to establish better bounds for the celebrated small ball problem for Gaussian processes. Equally important for the workshop is that many of the interesting spaces of functions occurring in numerical problems are well suited to the application of harmonic analysis.

As we understand better and better, these subjects are highly interrelated, and they are probably the most active and promising ones in the fields for the next decade. Bringing together a mix of junior and senior researchers from these diverse but interrelated subjects in a Dagstuhl seminar resulted in considerable progress both for the theory and the applications in these areas.

Seminars in applied mathematics and theoretical computer science typically consist of presentations, followed by short discussions in the plenum, and numerous informal discussions in smaller groups. In this seminar, we added another new feature. A moderator was assigned to three preselected talks (based on their particular relevance and on the experience of the speaker) in order to inspire a longer, in-depth discussion in the plenum. The three speakers were Jan Víbyral, Erich Novak, and Martin Hutzenthaler. The talks were scheduled as the first talks on Tuesday, Wednesday and Thursday. It was indeed very inspiring to witness the long and deep discussions following these special talks. We feel that this format was successful and should be used also in other workshops and conferences of the community.

The work of the attendants was supported by a variety of funding agencies. This includes the Deutsche Forschungsgemeinschaft, the Austrian Science Fund, the National Science Foundation (USA), and the Australian Research Council.

As always, the excellent working conditions and friendly atmosphere provided by the Dagstuhl team have led to a rich exchange of ideas as well as a number of new collaborations. Selected papers related to this seminar will be published in a special issue of the Journal of Complexity.

6.50 Advances and Challenges in Protein-RNA Recognition, Regulation and Prediction

Organizers: Rolf Backofen, Yael Mandel-Gutfreund, Uwe Ohler, and Gabriele Varani

Seminar No. 19342

Date: August 18–23, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.8.49

Participants: Amir Argoetti, Rolf Backofen, Marina Chekulaeva, Jörg Fallmann, Jan Gorodkin, Florian Heyl, Eckhard Jankowsky, Hilal Kazan, Julian König, Markus Landthaler, Donny Licatalosi, Yael Mandel-Gutfreund, Irmtraud Meyer, Neelanjan Mukherjee, Uwe Ohler, Yaron Orenstein, Teresa Przytycka, Michal Rabani, Andres Ramos, Olivia Rissland, Alexander Sasse, Michael Sattler, Michelle Scott, Michael Uhl, Charles E. Vejnar, Katharina Zarnack, Jianyang Zeng

DNA is often described as the blueprint of life, since it encodes all the information necessary for an organism to develop and maintain its biological functions. Single blueprints for specific functions are stored inside DNA regions called genes. The primary product produced (also termed expressed) from genes is RNA, which can either become biologically active itself (non-coding RNA or ncRNA) or is further translated into proteins (messenger RNA or mRNA), which then executes the gene functions. Given the astonishing complexity of biological functions, it is not surprising that the regulation of gene expression itself is a highly complex matter. Proteins, RNA, and DNA all can interact with each other, forming regulatory networks in order to control the expression of genes. To elucidate these networks, experimental scientists rely more and more on high-throughput methods, producing vast amounts of raw data. Computational methods to analyze these huge datasets are therefore of highest demand. The main focus of this seminar lies on RNA-protein and RNA-RNA interactions. In particular, transcriptome-wide binding patterns of RNA-binding proteins (RBPs), their computational predictability, and the biological effects of binding are discussed. Moreover, the seminar dealt with topics like combinatorial RBP binding prediction, RBP binding kinetics, RNA-RNA interaction prediction, subcellular RNA imaging, and RBP binding site classification. Regarding the computational methodology, several newly developed deep learning methods are presented, e.g. for RBP binding site prediction. Taken together, the aim of the seminar is to bring experimental and computational scientists together for the aforementioned topics and to engage them in fruitful discussions in order to:

- present the current experimental and computational methodologies,
- understand their implications, strengths, and limitations from first-hand experience,
- and spark ideas for developing new computational and experimental methods and improving on existing ones.

6.51 Computational Proteomics

Organizers: Nuno Bandeira and Lennart Martens

Seminar No. 19351

Date: August 25–30, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.8.70

Participants: Nuno Bandeira, Harald Barsnes, Pedro Beltrao, Sebastian Böcker, Robert Chalkley, Lieven Clement, Frank Conlon, David Creasy, Bernard Delanghe, Eric Deutsch, Maarten Dhaenens, Joshua Elias, Michael Götze, Rebekah Gundry, Sicheng Hao, Nils Hoffmann, Michael Hoopmann, Lukas Käll, Michelle Kennedy, Benoît Kunath, Lennart Martens, Magnus Palmblad, Hannes Röst, Renee Salz, Birgit Schilling, Brian Searle, Natalia Sizochenko, Stefan Tenzer, Yves Vandenbrouck, Hans Vissers, Olga Vitek, Juan Antonio Vizcaino, Mathias Wilhelm, Bernd Wollscheid, Roman Zubarev

The Dagstuhl Seminar 19351 'Computational Proteomics' discussed several key challenges of facing the field of computational proteomics. The topics discussed were varied and wide-ranging, and radiated out from the four topics set out at the start.

These four topics were (i) personally identifiable proteomics data; (ii) unique computational challenges in data-independent analysis (DIA) approaches; (iii) computational approaches for cross-linking proteomics; and (iv) the visual design of proteomics data and results, to communicate more clearly to the broad life sciences community. A cross-cutting topic was introduced as well, which focused on proteotyping in clinical trials as it brings many of the previous challenges together, by asking the logical but complex question of how proteomics approaches, data, and associated computational methods and tools can become part of routine clinical trial data acquisition, monitoring and processing.

Based on these initial topics, breakout sessions were organized around proteomics data privacy, dealing with data from DIA approaches, how to best utilize computational approaches to use cross-linking for structural elucidation, and the importance of visualisation of proteomics data and results to engender excitement for the field's capabilities in the life sciences in general. However, these breakout sessions in turn inspired additional breakout sessions on associated topics.

The DIA and cross-linking breakouts both yielded the issue of ambiguity in identification as a cross-cutting topic that merited its own dedicated breakout session. A closley related breakout session, derived from the proteomics privacy and DIA sessions, centered on open modification searches, which are now becoming feasible in proteomics for the first time, but which are also prone to potentially crippling ambiguity issues while raising even more complex privacy issues. The visual design breakout explicitly identified multi-omics data integration as a direct offshoot of its discussions, which led to a dedicated breakout session on this topic as well. Another emerging breakout session concerned public data, which was triggered by both the DIA and cross-linking topics because of their shared need to disseminate their respective specialised data and results in a standardised, uniform, and well-structured manner. Finally, the cross-linking and DIA topics also led to a breakout session on ion mobility, as this technological advance was seen as a key aspect in the future of these technologies.

Each of these breakout sessions had exciting outcomes, and gave rise to future research ideas and collaborations. The proteomics privacy breakout concluded that the field is now ready to delve in more detail into the issues surrounding proteomics data privacy concerns, and that a white paper will be written that can be used to propose policy and to inform the community. The DIA breakout identified three such future tasks: (i) to develop a perspective manuscript that will discuss peptide-centric and spectrum-centric FDR, as well as the effects of shared evidence; (ii) to conduct an experiment for testing DDA versus DIA on the same sample to discover the sampling space for precursors and fragments; and (iii) to conduct a second experiment for understanding target/decoy scoring for different decoy generation models using both synthetic and predicted target/decoy peptides. The cross-linking breakout concluded that a cross-linked ribosomal protein complex should be used as a standardized dataset publicly available to the community, while a 'Minimum Information Requirements About a Cross Linking Experiment (MIRACLE)' was proposed to unify results from many crosslinking tools. The results will also be presented at the Symposium on Structural Proteomics in Göttingen in November 2019. The visual design breakout came up with many fine-grained conclusions, but also

with an overall design philosophy which centered on three levels of technical detail, depending on the audience: i) interfaces for deatiled data exploration for experienced consumers; ii) interfaces with minimal technical information, focusing on high-level data for the specific scientific question for novice consumers; and iii) interfaces with only relevant information for clinical decision making (e.g. short list of proteins significantly affected by the disease) for clinicians.

The five offshoot breakouts described above also came to conclusions, and the interested reader is referred to the corresponding abstracts for details.

Overall, the 2019 Dagstuhl Seminar on Computational Proteomics was extremely successful as a catalyst for careful yet original thinking about key challenegs in the field, and as a means to make progress by setting important, high impact goals to work on in close collaboration. Moreover, during the Seminar, several highly interesting topics for a future Dagstuhl Seminar on Computational Proteomics were proposed, showing that this active and inspired community has not yet run out of challenges, nor out of ideas and opportunities!

6.52 Computation in Low-Dimensional Geometry and Topology

Organizers: Maarten Löffler, Anna Lubiw, Saul Schleimer, and Erin Moriarty Wolf Chambers

Seminar No. 19352

Date: August 25–30, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.8.84

Participants: Elena Arseneva, Maike Buchin, Benjamin Burton, Hsien-Chih Chang, Arnaud de Mesmay, Vincent Despré, Linda Kleist, Boris Klemz, Francis Lazarus, Maarten Löffler, Anna Lubiw, Clément Maria, Tim Ophelders, Hugo Parlier, Saul Schleimer, Lena Schlipf, André Schulz, Eric Sedgwick, Rodrigo I. Silveira, Jonathan Spreer, Frank Staals, Stephan Tillmann, Ivor van der Hoog, Birgit Vogtenhuber, Carola Wenk, Erin Moriarty Wolf Chambers, Alexander Wolff

One-dimensional structures embedded in higher-dimensional spaces are ubiquitous in both the natural and artificial worlds: examples include DNA strands, migration paths, planetary orbits, rocket trajectories, robot motion planning, chip design, and many more. These are studied in different areas of mathematics and computer science, under many names: knots, curves, paths, traces, trajectories, graphs, and others. However, researchers in many areas are just beginning to apply algorithmic techniques to find efficient algorithms for these structures, especially when more fundamental mathematical results are required. Broad examples of such problems include:

- classical algorithms on trajectories like the Fréchet distance as a way to formalize a distance measure as a curve changes;
- morphing between two versions of a common graph, which again tracks a higher dimensional space that corresponds to movement of a one-dimensional object;
- drawing and manipulating objects in three-manifolds, such as graphs, curves, or surfaces; and
- perhaps the simplest problem posed (in different ways) in all these areas, "how does one draw and morph a nice curve on a nice surface?"

This seminar was the second in a series. In the first seminar, the goal was to identify connections and seed new research collaborations along the spectrum from knot theory and topology, through to computational topology and computational geometry, and all the way to graph drawing. After the success of the first seminar, the goal for this second round was to continue and extend prior work, in particular by focussing on how objects change over time.

The seminar began with three overview talks from researchers in different areas (trajectory analysis, algorithmic topology, and graph drawing) to motivate and introduce problems which would fit the theme of changes over time in the representations of low-dimensional objects in higher dimensional spaces. We then invited all participants to describe open problems (most of which were circulated in advance of the meeting) that fit with the topic of the workshop and could benefit from broad expertise. For the remainder of the workshop we split into small working groups each focussed on a particular open problem.

Throughout the workshop we used Coauthor, a tool for collaboration designed by Erik Demaine (MIT), to share progress and updates among all the working groups. This, together with twice-daily progress reports, allowed us to share ideas and expertise among all participants, which was very effective. Another advantage was that we had a record of the work accomplished when the workshop ended.

Below, we (the organizers) describe the main working group topics and how they connected to the overarching theme. The abstracts of talks in the seminar and preliminary results from the working groups are also outlined later in this report.

One group worked on open questions that were motivated by 3-manifolds. In particular, they considered lower bounds for deciding the complexity of a knot or link equivalence, with a goal of finding specific knots that require many simplification moves. Their work involved both designing smaller examples, as well as doing larger scale exhaustive search using the software tool Regina.

Another group considered representations of graphs and hypergraphs by touching polygons in 3-d. They were able to leverage the dual graph of the polyhedral complex in 3d, and make progress on classifying which types of graphs could (or could

not) be realized. Their problem was primarily combinatorial, but the techniques used included several interesting topological arguments about embeddings of manifolds into 3d or into 3-manifolds.

Several groups considered problems about flows or morphs of curves in various settings. One question centered on visualizing actual embedded homotopies on a given surface; there is considerable prior work on how to compute such homotopies between curves quickly, but it generally focuses on computing the complexity of the homotopy as opposed to the actual sequence of simplifications needed. The group looked more closely at this algorithm, and was able to outline a proof that in fact an extension of that algorithm would generate the actual homotopy, for a slightly higher time cost. A second group considered curves in the plane, and investigated options for computing a "nice" morph between them. As the question was more vague, the group did quite a bit of background investigation on prior work, and then discussed a new technique based on 3-manifolds and normal surface theory which might lead to a new family of morphs. A third group looked at the problem of preprocessing a given curve so that the Fréchet distance to any other query curve could be efficiently computed, and were able to obtain improved time bounds for several variants of the problem.

In summary, the workshop fostered a highly collaborative environment where combining the expertise and knowledge of researchers from different communities allowed us to solve problems of common interest across those communities. A major theme was how connected the various problems could be; often, a proof technique or piece of literature suggested by a member of a different community proved useful or insightful to a group working in a different domain.

6

6.53 Logic and Learning

Organizers: Michael Benedikt, Kristian Kersting, Phokion G. Kolaitis, and Daniel Neider

Seminar No. 19361

Date: September 1–6, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.9.1

Participants: Isolde Adler, Molham Aref, Vaishak Belle, Michael Benedikt, Ismail Ilkan Ceylan, Victor Dalmau, Luc De Raedt, Dana Fisman, James Freitag, Ivan Gavran, Martin Grohe, Barbara Hammer, Daniel Huang, Nils Jansen, Brendan Juba, Kristian Kersting, Sandra Kiefer, Angelika Kimmig, Phokion G. Kolaitis, Egor Kostylev, Paul Krogmeier, Luis C. Lamb, Carsten Lutz, Mateusz Malinowski, Henryk Michalewski, Adithya Murali, Sriraam Natarajan, Daniel Neider, Dan Olteanu, Ana Ozaki, Madhusudan Parthasarathy, Lucian Popa, Martin Ritzert, Xujie Si, Dan Suciu, Christian Szegedy, Balder Ten Cate, Josef Urban, Steffen van Bergerem, Guy Van den Broeck, Zsolt Zombori

Motivation

Logic and learning are central to Computer Science, and in particular to AI research and allied areas. Alan Turing envisioned, in his paper "Computing Machinery and Intelligence" [1], a combination of statistical (*ab initio*) machine learning and an "unemotional" symbolic language such as logic. However, currently, the interaction between research in logic and research in learning is far too limited; in fact, they are often perceived as being completely distinct or even opposing approaches.

While there has been interest in using machine learning methods within many application areas of logic, the investigation of these interactions has usually been carried out within the confines of a single problem area. We believe that an interaction involving a broader perspective is needed. It would be fruitful to look for common techniques in applying learning to logic-related tasks, which requires looking across a wide spectrum of applications. It is also important to consider the ways that logic and learning, deduction and induction, can work together.

Design of the Seminar

The main aim of this Dagstuhl Seminar was to address the above problems by bring researchers from the logic and learning communities together and to create bridges between the two fields via the exchange of ideas ranging between the (seemingly) polar possibilities of the injection of declarative methods in machine learning and the use and applications of learning technologies in logical contexts. This included creating an understanding of the work in different applications, an increased understanding of the formal connections between these applications, and the development of a more unified view of the current attempts to organically reconcile deductive and inductive approaches. In order to structure these explorations, the focal points of the seminar were the following three distinct strands of interaction between logic and learning:

1. *Machine Learning for Logic*, including the learning of logical artifacts, such as formulas, logic programs, database queries and integrity constraints, as well as the application of learning to tune deductive systems.
2. *Logic for Machine Learning*, including the role of logics in delineating the boundary between tractable and intractable learning problems, the construction of formalisms that allow learning systems to take advantage of specified logical rules, and the use of logic as a declarative framework for expressing machine learning constructs.
3. *Logic vs. Machine Learning*, including the study of problems that can be solved using either logic-based techniques or via machine learning, an exploration of the trade-offs between these techniques, and the development of benchmarks for comparing these methods.

Summary of seminar activities

The seminar was attended by 41 researchers across various communities including logic, databases, Inductive Logic Programming (ILP), formal verification, machine learning, deep learning, and theorem proving. The membership consisted of senior and junior researchers, including graduate students, post-doctoral researchers, and industry experts. The seminar was conducted through talks and breakout sessions, with breaks for discussion between the attendees. There were three long talks, 21 short talks, and three breakout sessions on the discussion of open problems in logic and learning.

The talks consisted of: (i) presentation of recent advances in research questions and methodologies relating to the motivations discussed above; (ii) surveys of the state of research on various problems requiring the combination of deductive and inductive

reasoning as well as methodologies developed to address fundamental hurdles in this space; (iii) new perspectives on the organic combination of logical formulations and methods with machine learning in specific application domains; (iv) theoretical formulations and results on problems in learning logical representations; (v) demonstrations of state-of-the art tools combining logic and learning for applications such as theorem proving or entity resolution; (vi) presentation of research on challenge problems for the field of AI and intelligent reasoning.

The breakout sessions were conducted in three continuing parts, each spanning one session. The first part involved all the participants in a discussion of the current (small and large) open problems in AI, challenge problems for the field of intelligent systems, and research questions about defining specific goals representing a successful combination of inductive and deductive reasoning. This involved a deliberation of what problems were relevant, which problems could be potentially related to or dependent upon each other, and various suggestions to formalise commonly desired research goals. This session resulted in the choice of three broad areas for further specific discussion: (i) Explainable AI (ii) Injecting symbolic knowledge or constraints into neural networks, and (iii) Learning of logical formulae (first-order logic) from satisfaction on structures in a differentiable manner. The second part consisted of parallel thematic sessions on these three areas. Each thematic session was conducted in the form of a round-table discussion and was led by one or two participants who championed the theme. The third session brought all the participants together again to conclude with a summary of the ideas exchanged during the parallel sessions.

Conclusion

We consider the seminar a success. There is a growing need to enable the disparate communities of logic and learning to interact with each other, and we noted from the seminar that researchers from each community appreciated the perspective offered by the other, often identified techniques used by the other community that could be imported into their own, and, interestingly, were in agreement about the relevant and important problems of the day. The format of the seminar including ample time for discussions and breakout sessions received positive feedback from the participants.

References

1 A. M. Turing, “Computing machinery and intelligence”, Mind, vol. LIX, pp. 433–460, October 1950

6

6.54 Deduction Beyond Satisfiability

Organizers: Carsten Fuhs, Philipp Rümmer, Renate Schmidt, and Cesare Tinelli

Seminar No. 19371

Date: September 8–13, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.9.23

Participants: Alexander Bentkamp, Nikolaj S. Bjørner, Maria Paola Bonacina, Florent Capelli, Warren Del-Pinto, Rayna Dimitrova, Pascal Fontaine, Florian Frohn, Carsten Fuhs, Jürgen Giesl, Alberto Griggio, Arie Gurfinkel, Reiner Hähnle, Matthias Heizmann, Benjamin Kaminski, Laura Kovács, Quang Loc Le, Alexander Leitsch, Anthony W. Lin, Joao Marques-Silva, David Monniaux, Alexander Nadel, Claudia Nalon, Naoki Nishida, Quoc Sang Phan, Ruzica Piskac, Albert Rubio, Philipp Rümmer, Andrey Rybalchenko, Renate Schmidt, Martina Seidl, Viorica Sofronie-Stokkermans, Sorin Stratulat, Andrzej Szalas, Tachio Terauchi, Cesare Tinelli, Sophie Tourret, Andrei Voronkov, Uwe Waldmann, Christoph Weidenbach, Thomas Wies, Sarah Winkler

This report contains the program and outcomes of Dagstuhl Seminar 19371 on "Deduction Beyond Satisfiability" held at Schloss Dagstuhl, Leibniz Center for Informatics, during September 10–15, 2017. It was the thirteenth in a series of Dagstuhl Deduction seminars held biennially since 1993.

Research in automated deduction has traditionally focused on solving decision problems, which are problems with a binary answer. Prominent examples include proving the unsatisfiability of a formula, proving that a formula follows logically from others, checking the consistency of an ontology, proving safety or termination properties of programs, and so on. However, automated deduction methods and tools are increasingly being used to address problems with more complex answers, for instance to generate programs from formal specifications, compute complexity bounds, or find optimal solutions to constraint satisfaction problems.

In some cases, the required extended functionality (e.g., to identify unsatisfiable cores) can be provided relatively easily from current deduction procedures. In other cases (e.g., for Craig interpolation, or to find optimal solutions of constraints), elaborate extensions of these procedures are needed. Sometimes, altogether different methods have to be devised (e.g., to count the number of models of a formula, compute the set of all consequences of an ontology, identify missing information in a knowledge base, transform and mine proofs, or analyze probabilistic systems). In all cases, the step from yes/no answers to such extended queries and complex output drastically widens the application domain of deductive machinery. This is proving to be a key enabler in a variety of areas such as formal methods (for software/hardware development) and knowledge representation and reasoning.

While promising progress has been made, many challenges remain. Extending automated deduction methods and tools to support these new functionalities is often intrinsically difficult, and challenging both in theory and implementation. The scarcity of interactions between the involved sub-communities represents another substantial impediment to further advances, which is unfortunate because these sub-communities often face similar problems and so could greatly benefit from the cross-fertilization of ideas and approaches. An additional challenge is the lack of common standards for interfacing tools supporting the extended queries. Developing common formalisms, possibly as extensions of current standard languages, could be as transformative to the field as the introduction of standards such as TPTP and SMT-LIB has been in the past.

This Dagstuhl seminar brought together researchers working on deduction methods and tools that go beyond satisfiability and other traditional decision problems; specialists that work on advanced techniques in deduction and automated reasoning such as model counting, quantifier elimination, interpolation, abduction, or optimization; and consumers of deduction technology who need answers to more complex queries than just yes/no questions.

The unifying theme of the seminar was how to harness and extend the power of automated deduction methods to solve a variety of non-decision problems with useful applications. Research questions addressed at the seminar were the following:

- What kind of information should be passed to a "beyond satisfiability" deduction tool, and what information should be returned to the user? The goal should be to enhance the understanding of related concepts in different subfields and applications, and to converge towards common formalisms.
- How can current ideas, results and systems in one sub-community of researchers and practitioners benefit the needs of other communities?
- What are outstanding challenges in using and building deduction tools to attack logical problems with complex answers?

6.55 Application-Oriented Computational Social Choice

Organizers: Umberto Grandi, Stefan Napel, Rolf Niedermeier, and Kristen Brent Venable

Seminar No. 19381

Date: September 15–20, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.9.45

Participants: Haris Aziz, Dorothea Baumeister, Abdelhak Bentaleb, Sylvain Bouveret, Florian Brandl, Felix Brandt, Robert Bredereck, Markus Brill, Jiehua Chen, Cristina Cornelio, Ronald de Haan, Edith Elkind, Ulle Endriss, Piotr Faliszewski, Joseph Godfrey, Umberto Grandi, Davide Grossi, Ayumi Igarashi, Christian Klamler, Sascha Kurz, Martin Lackner, Jérôme Lang, Annick Laruelle, Omer Lev, Andrea Loreggia, Nicola Frederike Maaser, Janelle C. Mason, Nicholas Mattei, Nicolas Maudet, Reshef Meir, Vincent Merlin, Stefan Napel, Rolf Niedermeier, Arianna Novaro, David Pennock, Dominik Peters, Elven Priour, Jörg Rothe, M. Remzi Sanver, Ehud Shapiro, Karishma Rajesh Sharma, Piotr Skowron, Arkadii Slinko, Pietro Speroni di Fenizio, Nimrod Talmon, Paolo Turrini, Kristen Brent Venable, Toby Walsh, Lirong Xia

Computational social choice (COMSOC) combines models from political science and economics with techniques from computer science, to analyze collective decision processes from a computational perspective. Classical contributions include the study of the computational barriers to various forms of manipulation in elections, the definition of novel procedures for distributed resources among a group of human or artificial agents, as well as the study of complex collective decisions such as multi-winner voting rules and voting in combinatorial domains. COMSOC is a thriving field of research, with an international bi-annual workshop now at its 7th edition and a handbook published in 2016 which structures more than a decade of research, but future success will depend on the practical applicability of its findings. The purpose of this seminar was to address this challenge by stimulating application-driven research in computational social choice, i.e., theoretical studies modeling existing practical problems in all their complexity.

Four areas of COMSOC, which have already proven or bear particular potential for synergies and applicability to real-life problems, were identified as the focus of the seminar. Each of these areas addresses present-day challenges that provide an opportunity for an interdisciplinary approach building on contributions from computer scientists, economists, mathematicians, and political scientists:

- **Recommender systems** is a very successful application that combines several artificial intelligence techniques. Indeed, there have been few other examples of autonomous reasoning tools with comparable impact and pervasiveness in practice.
- **Fair division** has already proven a successful testbed for the application of theoretical work, thanks for the recently launched Spliddit webpage, which provides a user-friendly implementation for a number of algorithms in this field. This experience poses a number of questions and challenges for application-oriented research in fair division and beyond, such as data collection and analysis, possibly leading to new theoretical problems.
- **Interactive democracy** comprises a variety of approaches to make democratic processes more engaging and responsive. For instance, successful design and implementation of online decision platforms presents a multidisciplinary research challenge.
- **Real electoral systems** often have features that are absent in the single or multi-winner systems analyzed in textbooks and scientific papers. Voting theory and computational methods can help to identify non-monotonicity problems of real electoral systems, to provide normative benchmarks for institutional design, and to conduct influence and performance comparisons of different voting arrangements.

The Dagstuhl Seminar 19381 "Application-Oriented Computational Social Choice" brought together 46 invited participants of 15 different nationalities from 4 different continents, with three additional participants choosing to attend our seminar before participating to the Heidelberg Laureate Forum. The list of participants included researchers in Computer Science, Economics, and Political Science, three researchers from the industry (Microsoft, IBM, WinSet Group), and a lab technician.

For each of the focus topics described above, a 1-hour survey was prepared by one of the participants, obtaining an up-to-date overview of current research in the field and its main open problems. Each survey was scheduled on a different day, with 26 regular talks by participants complementing them in the program. Two rump sessions at the beginning of the week allowed a number of the participants to present recent findings, open problems and on-going research in a quick and informal way, stimulating the discussion for the rest of the week.

Given the focus of the seminar on application-oriented research, a special session was dedicated to the presentation of

software developed by researchers participating to the seminar. Voting platforms were presented (Whale[54] and OPRA[55]), a library for preference data (Preflib[56]), a platform for online deliberation and consensus building (Vilfredo[57]), as well as a number of tools to support experimental research in social choice. Moreover, the seminar hosted three live voting experiments during the week, two of which used a mobile experimental laboratory that was brought to Dagstuhl thanks to French CNRS and the help of a lab technician from University of Rennes. A detailed report of the experiments and an abstract of all the talks can be found below.

At the beginning of the week short sessions were reserved for individual self-introductions and for the proposition of potential group work. The organisers chose not to organize groups in advance, but to let them form in an iterative fashion during the seminar. A number of proposals were first made, then discussed and adapted, before participants signed up for specific group sessions. A total of 6 hours during the week was dedicated to group works, which led to significant advancements – a detailed report can be read below.

Overall, judging both from anecdotal personal feedback as well as the official results from the anonymous "Survey for Dagstuhl Seminar 19381" (with a median score of 10 out of 11 on the summary question "All in all, how do you rate the scientific quality of the seminar?" and similarly positive answers on the mix of participants, working atmosphere, etc.), the seminar was a very successful experience. It stimulated an already thriving research field to explore more applied research topics and scout for real-world problems. It allowed researchers to get first hand experience on how to run voting experiments, either on an Internet voting platform or in a laboratory, and allowed them to share their research practices. The work conducted in the groups was overall fruitful, already resulting in some paper drafts under preparation. The few suggestions for improvements mostly related to further broadening the mix of participants (more PhD students and junior researchers, more colleagues from nearby fields) and having a slightly less dense program (shorter talks, more time for work in small groups or unplanned activities).

The organisers wish to thank all the Dagstuhl staff for their professional support, the participants of the seminar for their positive attitude and enthusiasm, and the two collectors for putting together the abstracts that compose this report.

[54] https://whale.imag.fr/
[55] https://opra.cs.rpi.edu/polls/main
[56] http://www.preflib.org/
[57] https://www.vilfredo.org/

6.56 Data Ecosystems: Sovereign Data Exchange among Organizations

Organizers: Cinzia Cappiello, Avigdor Gal, Matthias Jarke, and Jakob Rehof

Seminar No. 19391

Date: September 22–27, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.9.66

Participants: Cinzia Cappiello, Ugo de' Liguoro, Yuri Demchenko, Elena Demidova, Boris Düdder, Bernadette Farias Lóscio, Avigdor Gal, Sandra Geisler, Benjamin Heitmann, Fritz Henglein, Matthias Jarke, Jan Jürjens, Maurizio Lenzerini, Wolfgang Maaß, Paolo Missier, Boris Otto, Elda Paja, Barbara Pernici, Frank Piller, Andreas Rausch, Jakob Rehof, Simon Scerri, Julian Schütte, Egbert Jan Sol, Gerald Spindler, Maria-Esther Vidal

The design of *data ecosystems*, infrastructures for the secure and reliable data exchange among organizations, is considered as one of the key technological enablers for digitalization and the digital economy of the future. Several applied research initiatives and industry consortia provide substantive evidence of this trend e.g., the Industrial Internet Consortium (IIC)[58] formed in the USA, the Industrial Data Space (IDS) founded in Germany and the associated consortium International Data Space Association (IDSA)[59]. Most of these initiatives aim to provide a *reference architecture* for dealing with (i) *governance* aspects related to the definition of policies and conditions able to norm the participation to the data ecosystem, (ii) *security* aspects related to the definition of policies and infrastructures for guaranteeing a trusted and secure exchange of data, (iii) *data and service management* aspects related to representation models and exchange formats and protocols, and (iv) *software design* principles related to the realization of the architectural components and their interaction.

All these aspects have been discussed in the seminar and the main findings are described in this report. In addition, a central new aspect of data ecosystems that we considered in the seminar lies in the view of data as having an economic value next to its intrinsic value to support operational and decisional core business activities. This means that in the data ecosystem, data is typically considered both a business asset and a business commodity which may be priced and sold in some form (e.g., data provisioning service or raw data) according to contracts.

As testified by the amount and variety of problems described above, the creation of such ecosystems poses many challenges cutting across a wide range of technological and scientific specializations. For this reason, the seminar involved researchers from different communities. Interdisciplinary discussions gave the possibility to analyze different perspectives and to achieve valuable outcomes presented in this report, such as a wide set of research challenges and the definition of interesting use cases for the further development of data ecosystems. Details about the activities carried out during the seminar are provided in the following.

Overview of the activities

The seminar took place from Monday September 23 until Friday September 27. The seminar program encompassed four invited talks (keynotes and tutorials) on the first day (Sep. 23rd), by Gerald Spindler (law and ethics), Frank Piller (ecosystems and business models), Maurizio Lenzerini (data integration), and Boris Otto (International Data Space). After discussions related to the talks and tutorials, the remaining afternoon was spent structuring (through joint discussion) the coming days of the seminar and group structure. As a result, group structure was based on a thematic structure encompassing three groups, one for each of the topic areas Business, Data, and Systems. Tuesday Sept. 24 began with a breakout into groups and election of scribes in each of the three groups (Business, Data, and Systems), and the remainder of the day was taken up by parallel group sessions in the three groups. Wednesday Sept. 25 began with a joint session where each of the groups presented their work, which was then discussed jointly. The afternoon (until the

[58] https://www.iiconsortium.org/
[59] https://www.internationaldataspaces.org/

excursion) was taken up by joint discussion on report structure. The morning of Thursday Sept. 26 encompassed joint discussion on a proposed joint manifesto as well as group discussions on application domains and application scenarios (topic areas were Health, SmartCities, Industry 4.0). The afternoon was taken up by continued group discussions and ended with group presentations and joint discussion on application domains and application scenarios. There was also further discussion on report structure at the end of the day. The manifesto was subject to very lively discussion in the evening, after dinner. Friday Sept. 27, the last day of the seminar, was devoted to wrap-up (conclusions, summary, and report process) followed by joint discussion on relations between Systems, Data and Business views on the overall topic of the seminar.

The outcome of the seminar, which is documented in the remainder of this report, encompasses summaries of the group discussions and the joint manifesto.

Fig. 6.15
"Dagstuhl trip report: learning and teaching programming language semantics" Blog post by 19281 Dagstuhl Seminar participant Amy J. Ko. https://medium.com/bits-and-behavior/dagstuhl-trip-report-learning-and-teaching-programming-language-semantics-b8d8d9007380. Photo courtesy of Amy J. Ko.

6.57 Comparative Theory for Graph Polynomials

Organizers: Jo Ellis-Monaghan, Andrew Goodall, Iain Moffatt, and Kerri Morgan

Seminar No. 19401

Date: September 29– October 4, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.9.135

Participants: José Aliste-Prieto, Nantel Bergeron, Cornelius Brand, Animesh Chaturvedi, Carolyn Chun, Anna De Mier, Jo Ellis-Monaghan, Graham Farr, Alex Fink, Delia Garijo, Daniela Genova, Emeric Gioan, Chris Godsil, Andrew Goodall, Krystal Guo, Orli Herscovici, Hendrik Jan Hoogeboom, Benjamin Jones, Nataša Jonoska, Louis H. Kauffman, Martin Kochol, Thomas Krajewski, Joseph Kung, Sergei Lando, Bodo Lass, Johann A. Makowsky, Iain Moffatt, Kerri Morgan, Steven Noble, Marc Noy, Jaeseong Oh, James Oxley, Vsevolod Rakita, Elena V. Ravve, Guus Regts, Adrian Tanasa, Maya Thompson, Peter Tittmann, Lluis Vena Cros, William Whistler, José Zamora Ponce

This 5-day Seminar built on the previous Dagstuhl Seminar 16241 together with several intervening workshops on graph polynomials, particularly those associated with William Tutte's Centenary, to advance an emerging comparative theory for graph polynomials. Graph polynomials have played a key role in combinatorics and its applications, having effected breakthroughs in conceptual understanding and brought together different strands of scientific thought. For example, the characteristic and matching polynomials advanced graph-theoretical techniques in chemistry; and the Tutte polynomial married combinatorics and statistical physics, and helped resolve long-standing problems in knot theory. The area of graph polynomials is incredibly active, with new applications and new graph polynomials being discovered each year. However, the resulting plethora of techniques and results urgently requires synthesis. Beyond catalogues and classifications we need a comparative theory and unified approaches to streamline proofs and deepen understanding.

The Seminar provided a space for the cross-fertilization of ideas among researchers in graph theory, algebraic graph theory, topological graph theory, computational complexity, logic and finite model theory, and biocomputing and statistical mechanics applications. There is a long history in this area of results in one field leading to breakthroughs in another when techniques are transferred, and this workshop leveraged that paradigm. More critically, experts in the field have recently begun noticing strong resonances in both results and proof techniques among the various polynomials. The species and genera of graph polynomials are diverse, but there are strong interconnections: in this seminar we worked towards a general theory that brings them together under one family. The process of developing such a theory of graph polynomials exposes deeper connections, giving great impetus to both theory and applications. This has immense and exciting potential for all those fields of science where combinatorial information needs to be extracted and interpreted.

The seminar was roughly organized according to the following themes:

- **Unification:** General frameworks for graph polynomials including meta-problems, K-theory, Second Order Logic, and Hopf algebras.
- **Generalizations:** Polynomial invariants for graphs with added structure (e.g. digraphs, ribbon graphs) or more general "underlying" combinatorial structures (e.g. matroids, Δ-matroids).
- **Distinction:** Distinguishing power of graph invariants (equivalence and uniqueness up to isomorphism with respect to a given graph polynomial, interrelations among graph polynomials, properties of graph polynomials).
- **Applications:** Applications of graph polynomials in other disciplines (e.g. self-assembly, sequencing, quantum walks, statistical mechanics, knot theory, quantum Ising model).
- **Conjectures:** Breakthrough conjectures (outstanding open problems whose resolution would have a broad impact on the understanding of graph polynomials).
- **Complexity:** Computational complexity and computational methods.

6

6.58 Social Agents for Teamwork and Group Interactions

Organizers: Elisabeth André, Ana Paiva, Julie Shah, and Selma Šabanovic

Seminar No. 19411

Date: October 6–11, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.10.1

Participants: Patrícia Alves-Oliveira, Elisabeth André, Tony Belpaeme, Cindy Bethel, Merijn Bruijnes, Filipa Correia, Mary Ellen Foster, Kobi Gal, Hatice Gunes, Dirk Heylen, Sviatlana Höhn, Malte Jung, James Kennedy, Franziska Kirstein, Stefan Kopp, Iolanda Leite, Andrea Marrella, Samuel Mascarenhas, Christoforos Mavrogiannis, Mark Neerincx, Catharine Oertel, Michio Okada, Ana Paiva, Simon Parsons, Catherine Pelachaud, André Tiago Abelho Pereira, Rui Prada, Laurel Riek, Sarah Sebo, Julie Shah, Elaine Short, Elizabeth Sklar, Marynel Vázquez, Hannes Högni Vilhjálmsson

As artificial agents and social robots become more prominent in our lives, they will also increasingly become parts of the groups and teams in which people spend much of their time. The objective of this Dagstuhl Seminar was to explore and discuss theories, methods, and techniques for building embodied social agents (including robots) that can operate in groups as members of a mixed team consisting of humans and agents. Recent advances in AI, and particularly in conversational agents, are likely to lead to an increased placement of agents in groups, covering a variety of application scenarios including healthcare, education, the workplace, and the home. Platforms such as Amazon Echo, Google Home, and new social robots such as Nao, Pepper, and Aibo facilitate such placement. Studies with robots in open-ended environments, including homes and public spaces, also suggest that people often engage with robots in such contexts in groups, rather than just individually. Yet, existing research on human-agent interaction and human-robot interaction so far focuses mostly on one-on-one interactions between a human and a social agent. To stimulate growing research in settings where one or more humans interact with multiple agents or robots, this seminar focused on human-agent communication, interaction, and teamwork in groups. As such, we discussed how agents shape the dynamics of groups, how agents and robots are able to perceive other members of a group and how they relate to each other, and how to move from one-to-one interactions to multi-party interactions of agents and humans in groups and teams. By bringing together researchers from different communities, such as human-robot interaction, multi-agent systems, social psychology, and organizational studies, we aim to generate common ground and new approaches in this interdisciplinary area. While this new domain of inquiry relies on existing research at the intersection between AI, robotics, and the social sciences, our aim is to highlight open questions that current work has not sufficiently addressed.

6.59 Quantum Cryptanalysis

Organizers: Michele Mosca, Maria Naya-Plasencia, and Rainer Steinwandt

Seminar No. 19421

Date: October 13–18, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.10.47

Participants: Gorjan Alagic, Daniel C. Apon, Daniel J. Bernstein, Jean-François Biasse, Christian Bischof, Xavier Bonnetain, Harry Buhrman, Jintai Ding, Martin Ekera, Philippe Gaborit, András Gilyén, Maria Isabel González Vasco, Sean Hallgren, Akinori Hosoyamada, David Jao, Samuel E. Jaques, Stacey Jeffery, Antoine Joux, Elena Kirshanova, Thijs Laarhoven, Bradley Lackey, Tanja Lange, Alexander May, Shaun Miller, Dustin Moody, Michele Mosca, Priyanka Mukhopadhyay, Maria Naya-Plasencia, Phong Q. Nguyen, Ray Perlner, Edoardo Persichetti, Rachel Player, Thomas Pöppelmann, Yu Sasaki, John M. Schanck, André Schrottenloher, Nicolas Sendrier, Yixin Shen, Daniel C. Smith-Tone, Rainer Steinwandt, Adriana Suárez Corona, Jean-Pierre Tillich, Iggy van Hoof, Fernando Virdia, Thomas Wunderer, Bo-Yin Yang

Motivation and scope

This fifth installment of a Dagstuhl seminar on *Quantum Cryptanalysis* was heavily informed by NIST's ongoing standardization effort in post-quantum cryptography. Several NIST employees attended the seminar and lead a discussion session on the topic. As one would hope hoped for, many talks had an algorithmic focus. Two areas were of particular interest for this seminar:

Quantum cryptanalytic progress. Identifying new cryptanalytic improvements that make use of quantum algorithms and expanding the applicability of the best known cryptanalytic attacks by means of quantum technology. Different quantum attack models can be considered here, and attack models that are close to being realizable with today's technology are particularly relevant. We want to fully leverage quantum computing, including expected mid-term advancements.

Quantum resource estimation. Establishing reasonably precise quantum resource counts for cryptanalytic attacks against symmetric and asymmetric schemes, especially for problem instances and parameter choices that are actually deployed or considered for standardization for future deployment. In addition to logical resources, understanding the overhead caused by handling imperfections of quantum hardware is of interest.

In addition to original quantum cryptanalytic research, the program included presentations with a strong survey component, explaining key concepts of particular areas within post-quantum cryptography. Deviating from prior editions, this time we did not include a presentation to document the status of the development of quantum hardware. Such a talk could have been a welcome addition, but the seminar program was already packed with a substantial number of relevant cryptanalytic results, and it was important to leave sufficient time for discussions.

Organization

Following the organization of the prior quantum cryptanalysis seminars in Dagstuhl, for this fifth edition, again experts from academia, government, and industry came together. We re-invited a number of leading experts in the field from the prior quantum cryptanalysis seminar edition, and at the same time invited several new participants. This included in particular young scientists, who entered this exciting research area more recently. In total, we had with 46 participants a slightly larger number of participants than in the preceding meeting. In line with the Dagstuhl tradition and with prior quantum cryptanalysis seminars, for Wednesday afternoon we left the schedule open. Seminar participants could devote the afternoon to an excursion, to discussions, or to work on their research.

Results and next steps

At this point, communication and collaboration between the classical cryptographic and the quantum algorithmic research communities has become very fruitful, and it seems fair to say that this seminar is also of significant value in supporting ongoing standardization efforts in post-quantum cryptography. In addition to quantum cryptanalytic results on asymmetric cryptography, more results on symmetric cryptography are emerging. There is still substantial research potential – and research need – in quantifying security margins in the presence of quantum computing, and the field keeps moving fast. Improved software tools become available to analyze quantum resources and describe quantum algorithms, bringing research in quantum cryptanalysis closer together with areas in traditional computer science.

6

6.60 Theory of Randomized Optimization Heuristics

Organizers: Carola Doerr, Carlos M. Fonseca, Tobias Friedrich, and Xin Yao

Seminar No. 19431

Date: October 20–25, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.10.61

Participants: Youhei Akimoto, Denis Antipov, Anne Auger, Thomas Bäck, Thomas Bartz-Beielstein, Hans-Georg Beyer, Vivek Shripad Borkar, Dimo Brockhoff, Maxim Buzdalov, Arina Buzdalova, Francisco Chicano, Alexandre Chotard, Benjamin Doerr, Carola Doerr, Anton V. Eremeev, Carlos M. Fonseca, Tobias Friedrich, Tobias Glasmachers, Nikolaus Hansen, Thomas Jansen, Timo Kötzing, Martin S. Krejca, Per Kristian Lehre, Johannes Lengler, Andrei Lissovoi, Manuel López-Ibánez, Rolf H. Möhring, Frank Neumann, Pietro S. Oliveto, Luc Pronzato, Jonathan E. Rowe, Günter Rudolph, Ofer M. Shir, Patrick Spettel, Dirk Sudholt, Andrew M. Sutton, Olivier Teytaud, Dirk Thierens, Vida Vukašinovic, Markus Wagner, Elizabeth Wanner, Thomas Weise, Carsten Witt, Xin Yao, Christine Zarges, Anatoly Zhigljavsky

Efficient optimization techniques affect our personal, industrial, and academic environments through the supply of well-designed processes that enable a best-possible use of our limited resources. Despite significant research efforts, most real-world problems remain too complex to admit exact analytical or computational solutions. Therefore, heuristic approaches that trade the accuracy of a solution for a simple algorithmic structure, fast running times, or an otherwise efficient use of computational resources are required. Randomized optimization heuristics form a highly successful and thus frequently applied class of such problem solvers. Among the best-known representatives of this class are stochastic local search methods, Monte Carlo techniques, genetic and evolutionary algorithms, and swarm intelligence techniques.

The theory of randomized optimization heuristics strives to set heuristic approaches on firm ground by providing a sound mathematical foundation for this important class of algorithms. Key challenges in this research area comprise optimization under uncertainty, parameter selection (most randomized optimization heuristics are parametrized), the role and usefulness of so-called *crossover* operations (i.e., the idea of creating high-quality solution candidates by recombining previously evaluated ones) and, more generally, performance guarantees for advanced heuristics such as population-based techniques, estimation-of-distribution algorithms, differential evolution, and others.

Dagstuhl Seminar 19431 on "Theory of Randomized Optimization Heuristics" was a continuation of the seminar series originally on "Theory of Evolutionary Algorithms". Today the field extends far beyond evolutionary algorithms – a development that previous Dagstuhl seminars have significantly influenced.

While the previous seminar 17191 had a very strong focus on methodological questions and techniques needed to analyze stochastic optimization heuristics, the present seminar had among its three main focus topics chosen to foster interaction with two strongly linked research communities that were not previously represented in the seminar series: stochastic control theory and empirical benchmarking of randomized optimization heuristics.

Recent work has shown that there is a very close link between the theory of randomized optimization heuristics and stochastic control theory, both regarding the nature of the "systems" of interest and the analytical techniques that have been developed in the two communities. At the seminar, we have explored these affinities through the two invited presentations of Luc Pronzato and Vivek Borkar, through contributed talks highlighting different aspects studied in both communities (e.g., the presentation on one-shot optimization by Olivier Teytaud), and through focussed breakout sessions, in particular the one fully dedicated to *Connection between the analysis of evolution strategies and estimation of distribution algorithms and the analysis of stochastic approximation and ordinary differential equations*, in which interesting similarities and differences between the two fields were identified.

The second focus topic of Dagstuhl Seminar 19431 was benchmarking of optimization heuristics. Benchmarking plays a central role in empirical performance assessment. However, it can also be an essential tool for theoreticians to develop their mathematically-derived ideas into practical algorithms, thereby encouraging a principled discussion between empirically-driven and theoretically-driven researchers. Benchmarking has been a central topic in several breakout sessions, for example those on *Competitions and Benchmarking*, *Algorithm Selection and Configuration*, but also the breakout session on *Multi-Objective Optimization*. A survey of best practices in empirical benchmarking has been kick-started in the breakout session on *Benchmarking: Best Practices and Open Issues*.

Discussing the mathematical challenges arising in the performance analysis of randomized heuristics has always been a central topic in this Dagstuhl seminar series. Among other achievements, important connections between continuous and

discrete optimization have been established, most notably in the form of drift theorems, which are typically applicable regardless of the nature of the search space. Apart from such methodological advances, we have also observed two other trends bridging discrete and continuous optimization: (i) an increased interest in analyzing parameter-dependent performance guarantees, and (ii) the recent advances in the study of estimation of distribution algorithms, which borrow techniques from both discrete and continuous optimization theory. These topics have been discussed in the invited talk of Youhei Akimoto, in several contributed presentations, and in the breakout sessions on *Measuring Optimization Progress in an Invariant Way for Comparison-Based Algorithms* and on *Mixed-Integer Optimization*.

Apart from these focus topics, we have discussed a large number of different aspects related to the theoretical analysis of optimization heuristics, including brainstorming sessions on doing "good" research, organizing a repository to share lecture materials, and discussing the role of uncertainty in heuristic optimization, the connections between experimental design and one-shot optimization, the importance of neutral representations, and differences between stochastic gradient descent methods and evolution strategies, to give but a few examples.

Organization

The seminar hosted the following type of events:

- Five invited talks of 30 minutes each:
  - Youhei Akimoto on *Expected Runtime Bound for the (1+1)-Evolution Strategy*
  - Vivek Borkar on *Overview of Stochastic Approximation and Related Schemes*
  - Pietro S. Oliveto on *What is Hot in Evolutionary Computation (Part 2)*
  - Luc Pronzato on *Dynamical Search*
  - Carsten Witt on *What is Hot in Evolutionary Computation (Part 1)*
- 20 contributed talks of around 15-20 minutes
- Four "flash talks" of about 10 minutes
- Eleven parallel breakout sessions in various different formats, ranging from brainstorming on the purpose and future of theory research through actual problem solving on one-shot optimization to kick-starting a survey on best practices on benchmarking optimization heuristics.

All presentations were plenary, i.e., in a single session, while the breakouts were organized in parallel working groups, to allow for focused and specialized discussions. As in previous years, the breakout sessions were very well perceived, and can be considered a well-established format of this seminar series. As a result of these discussions, we are planning a workshop and a survey on benchmarking best practices. Several open problems have been proposed and discussed at the seminar, and we are confident that the seminar has helped to establish new collaborations.

Our traditional hike on Wednesday was a good opportunity to discuss in a less formal setting and to get to know each other. On Thursday evening, we had the special opportunity to hear Jonathan Rowe present activities of the Alain Turing Institute https://www.turing.ac.uk/, where he serves as Programme Director for Data Science for Science. Last, but not least, the wine-and-cheese party complemented the scientific activities with a relaxed social event.

We would like to thank the Dagstuhl team and all participants for making seminar 19431 a great success and a great pleasure to organize.

Carola Doerr (Sorbonne University – Paris, FR)
Carlos M. Fonseca (University of Coimbra, PT)
Tobias Friedrich (Hasso Plattner Institute – Potsdam, DE)
Xin Yao (Southern University of Science and Technology – Shenzen, CN)

6.61 Analysis of Autonomous Mobile Collectives in Complex Physical Environments

Organizers: Mario Gleirscher, Anne E. Haxthausen, Martin Leucker, and Sven Linker

Seminar No. 19432

Date: October 20–23, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.10.95

Participants: Étienne André, Andreas Bauer, Sergiy Bogomolov, Jörg Brauer, Radu Calinescu, Alessandro Fantechi, Frederik Foldager, Sibylle Fröschle, Mario Gleirscher, Anne E. Haxthausen, Christian Heinzemann, Jean-Baptiste Jeannin, Kim Guldstrand Larsen, Peter Gorm Larsen, Martin Leucker, Sven Linker, Stefan Mitsch, Laura Nenzi, Peter Csaba ölveczky, David Parker, Pedro Ribeiro, Masaki Waga

Motivation

Autonomous vehicles (AVs) are facing strong proof obligations. Individual AVs can be part of a *collective* (e.g. a platoon of utility vehicles on a farm field, a truck convoy on a highway, a convoy of passenger vehicles on urban road, an in-door aerial platoon, a railway convoy) and act within a heterogeneous environment of other collectives, for example, pedestrians, bicyclists, and motorcyclists. Multiple AVs might have to correctly and reliably negotiate their order of passing a crossing or reliably and robustly arrange in a certain work layout on agricultural land. Individuals and collectives in such environments, whether controlled in a centralised or distributed way, are subjected to change, uncertainty, and defects. Moreover, *complex environments* typically deny a comprehensive segregation of physical space and, hence, involve interactions with entities out of control (e.g. human-controlled machines, pedestrians, animals) and mostly also out of sight of an individual machine's (short-range) sensors.

Objective

This seminar was centred around an application challenge, the **Smart Farm**. Participants were encouraged to discuss how their research addresses typical **engineering tasks** (ETs; upper layer in Fig. 6.16) to be accomplished for the given challenge or for similar challenges. These tasks include

1. the identification, modelling, and analysis of operational situations in complex environments
2. real-time coordination, composition, and reconfiguration of machine collectives with a focus on (i) interaction with human-operated systems, humans, animals, infrastructure and (ii) situation-specific centralised or distributed control regimes
3. the determination of strongest safety and performance guarantees with a focus on (i) the estimation of upper resilience bounds of machine collectives and lower reliability bounds of individual machines and (ii) the determination of strongest guarantees under partial state knowledge, with minimal infrastructural support, and under reduced controllability.

In the discussions of how the ETs can be accomplished best, we also aimed at investigating **abstractions of defects and uncertainties**, for example:

- controller, communication, and infrastructure failures (e.g. erroneous vehicle-to-X connection and communication, deficient road infrastructure),
- undesired interference or disturbance of autonomous operation (e.g. malicious and unintended misuse; controller, communication, and infrastructure attacks),
- practical sensor uncertainties, actuator perturbations, and partial state knowledge.

Defects and uncertainties are crucial for constructing *realistic models* of the behavioural spectrum of mobile collectives and yet abstract enough to perform *practical reasoning*. Likewise, such models allow the necessary freedom to express ideal and actual behaviour, independent of whether such behaviour is desirable. This freedom can involve the use of non-deterministic models. In any case, a *(property) specification* would label some of the observable behaviours as *desirable*, some as *undesirable*, others somewhere in between (cf. quantitative verification). The more complete and precise such a specification, the better the distinction between correct, undesirable, and other classes of behaviours of a collective.

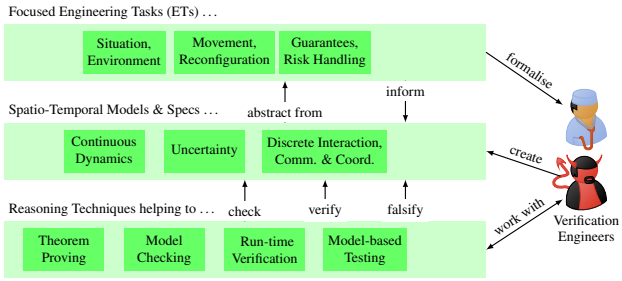

Fig. 6.16
Topic structure of the seminar.

Our **overall objective** with this seminar was to gain a common understanding of acceptable safety and performance of autonomous mobile collectives in presence of defects and other uncertainties typically occurring in complex open environments. The **overarching approach** of all seminar contributions was the **formal analysis and verification of behavioural correctness** under these assumptions (lower layer in Fig. 6.16) by using techniques such as, e.g. theorem proving, model checking, run-time verification, and model-based testing.

Our **central assumption for this seminar** was that the given application challenge or any similar challenges render individual methods for the analysis and verification of such systems insufficient. For example, in control-theoretic models such collectives are modelled by differential equations. Interaction within and among collectives and with their environment, governing these equations, cannot be easily encoded. Approaches that express such interactions well, however, typically struggle with the detailed description of the physical laws the AVs need to adhere to. Hence, for ensuring correct behaviour in such a setting, layered abstractions, corresponding models, and **specialised reasoning techniques have to be combined**.

Organisation

Before the seminar, we provided each participant with material about the *application challenge* (see Section 4.1 of the full report) together with list of *engineering tasks* and *research questions*. We encouraged the participants to apply their approach, if available, to at least one of the ETs of the application challenge and to answer at least one of the research questions. Alternatively, participants were invited to present any research and practical experiences related to the seminar topic and the challenge. Everyone was given the opportunity to give a full-length talk. Table 6.3 shows the seminar structure, the talks, and further sessions. After the welcome session, participants introduced themselves to the group. The rest of the seminar was organised into *talk sessions* and *break-out sessions*.

Table 6.3
Seminar schedule.

	Monday	Tuesday	Wednesday
9:00	Introductions	***Industry Challenges*** J. Brauer: *Verification of Autonomous Transport Systems - Some Industrial Prospects*	*Break-out session*
9:30		S. Fröschle: *Trustworthy identity and key management for mobile systems in transportation*	
10:00–10:30	break	break	break
10:30	***Individual Properties*** P.G. Larsen/F. Foldager: *A Journey Towards a Fleet of Autonomous Robots for Agricultural Field Operations*	***Uncertainty Modelling*** K.G. Larsen: *Synthesis of Safe, Optimal and Small Strategies for Advanced Driver Assistance using UPPAAL Stratego*	*Break-out and discussion*
10:50	J.B. Jeannin: *Collision avoidance and path replanning of individual farm robots*	D. Parker: *Probabilistic model checking for safety and performance guarantees*	*Closing discussion*
11:10	A. Fantechi: *Safety aspects of autonomous systems*	R. Calinescu: *Stochastic modelling underpinning the engineering of trustworthy autonomous systems*	
11:30	P.C. Ölveczky: *Formal modeling and analysis of real-time systems using Real-Time Maude*	M. Gleirscher: *Risk Structures*	
12:15–13:30	lunch	lunch	lunch
13:30	***Collective Properties*** M. Waga: *Optimization of the watering schedule by run-time and design-time analysis*	***Individual Properties*** C. Heinzemann: *Context Analysis and Requirements Derivation with SCODE*	
13:50	É. André: *White-box and black-box quantitative verification of timing properties*	S. Bogomolov: *Trusted Autonomous Systems: Verification Meets Falsification*	
14:10	P. Ribeiro: *Modelling and Verification using RoboChart*	S. Mitsch: *Modular Verification of Cyber-Physical Systems in KeYmaeraX*	
14:30	(spare)	(spare)	
15:00–15:30	break	break	
15:30			
16:00	*Break-out session*	*Break-out session*	
16:30			
17:00	*Discussion of results*	*Discussion of results*	
18:00	dinner	dinner	

Talks. In the talk sessions, we investigated several **research questions** from different angles. We had talks about (1) industry challenges, (2) the analysis and verification of properties of individual autonomous vehicles (two sessions), (3) the analysis and verification of properties of autonomous collectives, and (4) the modelling of uncertainty for the (quantitative) property verification of critical autonomous systems. Nine talks dealt with an **introduction of a specific verification approach** suitable for tackling an aspect of the application challenge, including a summary of the state-of-the-art of this approach. Four talks were about **industrial examples** of a nature similar to the Smart Farm, highlighting technical challenges, encountered issues, and perceived practical obstacles. Five talks focused on the **application of a particular approach to a particular aspect of the Smart Farm**, addressing some of the research questions.

In the following, we list the main questions and the participants whose talks highlighted a particular aspect of the corresponding question. For more details, see the list of talk abstracts below.

1. How can each ET be solved? How can we achieve safety in presence of distribution, mobility, and uncertainty? Which mechanisms fit best to ensure safety in the *application challenge*?
   Frederik Foldager and Peter Gorm Larsen
2. How do we model the systems and verify *safety and progress* properties? Can we always find acceptable Pareto optima over safety and performance, at traffic level, at the level of a collective, and for individual machines?
   Étienne André, Sergiy Bogomolov, Kim Larsen, David Parker
3. How can we *exploit the structure* of practical AVs and collectives to craft specific verification techniques (e.g. prevent state space explosion, identify fundamental theorems)?
   Stefan Mitsch, Pedro Ribeiro
4. Which benefits do we gain from *integrating* design-time verification, model-based testing, and run-time verification?
   Mario Gleirscher, Masaki Waga
5. How can verification techniques be incorporated into the *development process* of AVs?
   Jörg Brauer, Radu Calinescu, Alessandro Fantechi, Peter Csaba Ölvecky
6. Which *complications* arise from the verification of AVs and how can we mitigate the impact of these complications, particularly, during practical verification?
   Sibylle Fröschle, Christian Heinzemann

Break-Out Sessions. To stimulate interaction, we created *break-out groups* on each seminar day and on the following topics: challenges of verifying autonomous collectives, the challenge of uncertainty (using, e.g. quantitative verification, parametric model checking), abstractions of space & uncertainty, the impact of IT security issues on AV safety, and safe platooning. Additionally, several smaller groups (sometimes consisting of only two participants) met to discuss combinations and extensions of the topics they presented in their respective talks.

One break-out group focused on creating a big picture of the **challenges of verifying autonomous mobile collectives** in the Smart Farm. The identified problems include

- estimation of behavioural properties (e.g. exact arrival times of agents, dead-lock freedom of the plan), real-time interleaving of sensing and control, and finding the "sweet spot" between precision and performance when used at run-time,
- model checking at scale, when to use online or offline analysis for verification and synthesis (e.g. synthesis of distributed safety controllers for automatic repair/fallback),
- useful architectural abstractions, compositionality, and refinement (e.g. how to safely partition the tasks of a mission between system components or whole robots?),
- security of communication and robustness of control to communication glitches (e.g. how to integrate a jamming model into overall system verification?),
- languages/models for dealing with system failures (e.g. how to cope with failures of individual autonomous vehicles in the context of a collective?) and component failures (e.g. how to safely integrate machine learning into autonomous systems?), and
- safety in the presence of uncertainty (e.g. how to quantify uncertainty?, how to deal with uncertainty in parameters and in the structure of the system and the environment?).

Another group investigated **the challenge of uncertainty in modelling**, discussing how uncertainty (e.g. due to partial observability) can be dealt with in automated verification and how techniques such as quantitative verification can be used to solve verification problems with uncertainties in the considered parameters. Depending on the Smart Farm aspect to be tackled, state-of-the-art approaches include the use of interval abstractions for parameters, the calculation of confidence intervals for verification results, and the use of counterexample-guided abstraction refinement.

The break-out session on **space and uncertainty** stretched over all three days, and was concerned with the possible ways to specify spatial aspects, as well as how to incorporate uncertainty into such specifications. Our discussion proceeded on different topics. We discussed, which types of sensors allow robotic systems to gain spatial knowledge, and what levels of uncertainty can be expected. Based on this, we examined whether several layers of space are necessary and beneficial to specify both the systems and their desired properties (e.g., a discrete layer for planning high-level actions and a continuous layer, on which more local properties are ensured by controllers, as for example obstacle avoidance). Furthermore, we compared the different types of uncertainty, the level of spatial layers they occur on, and their impact on systems in the Smart Farm. This included a discussion of how much knowledge needs to be globally available, and what can be kept locally at the level of each individual entity. We realised that while the modelling scenario allowed for different levels of space and uncertainty, it was not easy and straightforward to identify necessary and interesting spatial properties to analyse. Hence, we agreed that the case study needs to allow for more degrees of freedom (e.g., different routes to reach physical targets, to permit several alternative plans).

The session about **IT security of farm collectives** focused on the aspect of communication security. First, the group identified the typical communication requirements between the actors of a smart farm such as: between a robot and a supervisory control (perhaps including a drone), between two robots that carry out a task on the same field (e.g. to carry out the task cooperatively or for collision avoidance), between a sensor and a control centre (e.g. for watering). Altogether, it became clear that the operation of a smart farm critically depends on the secure and timely communication between the various actors. It is also clear that in the setting of the smart farm the actors must communicate over wireless channels. Hence, the usual threats against communication over an open medium apply, e.g. message spoofing and manipulation, eavesdropping and jamming. On the one hand, this requires us to employ appropriate security protocols and key management, which can guarantee origin and message authenticity as well as confidentiality. On the other hand, this requires further measures against availability attacks such as jamming. The group focused on the threat of jamming.

While jamming cannot be prevented in an open system the general idea was to take a 'detect and mitigate' approach. For example, jamming can be detected by the absence of regular 'heartbeat' signals and by combination with visual channels. Mitigation strategies involve raising an alarm and removing the jamming device in a timely fashion while ensuring the system is not overly susceptible to false positives and denial-of-service attacks. Neither detection nor mitigation seemed trivial when discussed in detail. On the positive side, the verification methods and tools presented at the seminar could be used to evaluate possible strategies, and perhaps, even to synthesise them. Later on the group joined the break-out group on platooning, where communication is particularly critical.

In the break-out session on **safe platooning on the farm**, we discussed

1. the handling of *planned events* being part of the normal operation of a platoon (e.g. several farm vehicles, lorries and harvesters, form a platoon including leader election; a lorry wants to join or leave a harvesting platoon; a platoon with two consecutive lorries needs to be rearranged; a lorry decides to leave the platoon) and
2. the detection of *critical (not necessarily undesired) events* to be dealt with or to recover from during normal operation (e.g. a foreign vehicle, a farmer's car, enters the platoon area; communication error because of a jamming attack or a hardware failure disturbs the platoon controller; the current leader looses trustworthiness, e.g. because of being hacked, by deviating from the common goal of the platoon; farm workers enter the working area of the platoon).

Our discussions lead to a deeper understanding of the intricacies, both from the perspectives of different verification approaches and from the viewpoint of certification obligations. The results of our discussion are suitable for the identification of *formal properties* to be used as proof obligations in certification activities as well as the modelling of so-called protocol automata *describing the inter- and intra-modal behaviour* required to handle some of the mentioned events. Such models can then serve as a basis for hazard and risk assessment activities as well as for safety verification.

Outcomes and Conclusions

Our expectations for this first seminar were modest. We wanted to learn from each others' perspectives, to discuss available approaches, and to identify the hardest and most relevant **open challenges**.

Our discussions opened **paths to an integration and application of the presented theories and models** (middle layer in Fig. 6.16), particularly, continuous models (e.g. timed and hybrid automata), uncertainty models (e.g. Markov chains, probabilistic automata), communication and coordination models (e.g. timed process algebra). We investigated the use of such models in the context of various reasoning techniques (e.g.theorem proving, model checking, run-time verification, model-based testing). These discussions lay a basis for the *derivation of guidelines* on how the approaches, when applied to systems such as the Smart Farm, can be combined and/or enhanced to tackle the identified problems *in practical contexts subject to certification efforts*.

The attendees were from various fields such as formal verification, testing, certification, mechanical and control engineering, and embedded IT security, working at universities, in industry-oriented research institutes, or directly in industry. In this setting, we were able to **share experiences and insights from various application domains** (e.g. smart farming, smart energy systems, train/railway systems, automotive and transportation), to discuss issues of the Smart Farm scenario, and to examine potential research directions. Particularly, we observe that commonalities among the used approaches give rise to an *integrated and more versatile approach*. Our participants from industry receive the opportunity to convert any of these insights into lasting process improvements in their safety-critical domains. We expect our findings to be *relevant to regulatory authorities* in these domains.

In overall, we believe this seminar was an important step to **foster collaboration** of researchers and practitioners experienced with the different models and reasoning techniques, and to **initiate a research community** focusing on autonomous collectives of similar or even higher complexity than the Smart Farm. To that end, we are planning **further meetings** of the seminar's participants in the near future, to allow for further refinement of the models, and combinations of the methods presented. Additionally, we will further improve and **extend the modelling scenario**, so that a particular combination of specification and verification approaches can be explored in more detail. Eventually, we intend to collect our findings possibly in a special issue of a suitable journal.

Funding and Acknowledgements. Sven Linker was supported by the Engineering and Physical Sciences Research Council programme grant EP/N007565/1 (S4: Science of Sensor Systems Software). Mario Gleirscher was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under the Grant no. 381212925. We are grateful to Sibylle Fröschle for summarising the results of the IT security session. Further thanks go to Frederik Foldager for collecting and compiling the abstracts. We would like to spend sincere gratitude to all participants for their contributions and for their support and active engagement in making this seminar an insightful experience.

6.62 Programming Languages for Distributed Systems and Distributed Data Management

Organizers: Carla Ferreira, Philipp Haller, and Guido Salvaneschi

Seminar No. 19442

Date: October 27–31, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.10.117

Participants: Rohan Achar, Carlos Baquero, Annette Bieniusa, Uwe Breitenbücher, Sebastian Burckhardt, Surajit Chaudhuri, Natalia Chechina, Amit K. Chopra, Schahram Dustdar, Patrick Thomas Eugster, Carla Ferreira, Torsten Grust, Philipp Haller, Edward A. Lee, Heather Miller, Aleksandar Prokopec, Laurent Prosperi, Guido Salvaneschi, Manuel Serrano, Marc Shapiro, Marjan Sirjani, Peter Van Roy, Nobuko Yoshida, Damien Zufferey

Developing distributed systems is a well-known, decades-old problem in computer science. Despite significant research effort dedicated to this area, programming distributed systems remains challenging. The issues of consistency, concurrency, fault tolerance, as well as (asynchronous) remote communication among heterogeneous platforms naturally show up in this class of systems, creating a demand for proper language abstractions that enable developers to tackle such challenges.

Over the last years, language abstractions have been a key for achieving the properties above in many industrially successful distributed systems. For example, MapReduce takes advantage of purity to parallelize task processing; complex event processing adopts declarative programming to express sophisticated event correlations; and Spark leverages functional programming for efficient fault recovery via lineage. In parallel, there have been notable advances in research on programming languages for distributed systems, such as conflict-free replicated data types, distributed information-flow security, language support for safe distribution of computations, as well as programming frameworks for mixed IoT/cloud development.

However, the researchers that have been carrying out these efforts are scattered across different communities which include programming language design, type systems and theory, database systems and database theory, distributed systems, systems programming, data-centric programming, and web application development. This Dagstuhl Seminar brought together researchers from these different communities.

The seminar focused on answering the following major questions:

- Which abstractions are required in emergent fields of distributed systems, such as mixed cloud/edge computing and IoT?
- How can language abstractions be designed in a way that they provide a high-level interface to programmers and still allow fine-grained tuning of low-level properties when needed, possibly in a gradual way?
- Which compilation pipeline (e.g., which intermediate representation) is needed to address the (e.g., optimization) issues of distributed systems?
- Which research issues must be solved to provide tools (e.g., debuggers, profilers) that are needed to support languages that target distributed systems?
- Which security and privacy issues come up in the context of programming languages for distributed systems and how can they be addressed?
- What benchmarks can be defined to compare language implementations for distributed systems?

The seminar accomplished the goal of bringing together the research communities of databases, distributed systems, and programming languages. The list of participants includes 24 academic and industrial researchers from Austria, Belgium, France, Germany, Portugal, Sweden, Switzerland, UK, and USA, with complementary expertise and research interests. The group had a balanced number of senior researchers and junior researchers, as well as a strong industrial representation.

The scientific program comprised 28 sessions. The sessions devoted to individual presentations included 16 short talks with a maximum duration of 15 minutes and 6 long contributed talks with a maximum duration of 35 minutes. In addition, the seminar included 2 plenary sessions and 4 group sessions. The first two days of the seminar were dedicated to research talks, but it was ensured that each talk had allocated time for discussions and exchange of ideas. In the two following mornings there

were 3 plenary sessions and 2 parallel group sessions. The topics for these sessions were proposed and selected after a lively discussion between participants, where the most popular sessions were promoted to plenary and the remaining occurred in two parallel sessions. The scientific sessions discussed and collected open questions on the topics of: programming models and abstractions; security and privacy; static guarantees, type systems, verification; distributed computing for the edge; time, synchrony, and consistency; and persistency and serialization. There was also a social topic discussing further actions to bring the three communities together. Even though there are overlapping research interests, there is a difference of values between communities that needs to be acknowledged and tackled. Participants agreed on the goal of organizing follow-up events to further strengthen the connection among the database, the distributed systems and the programming languages communities. In particular, the importance of extending future events to Ph.D. students, for instance with an integrated Summer School, has been discussed.

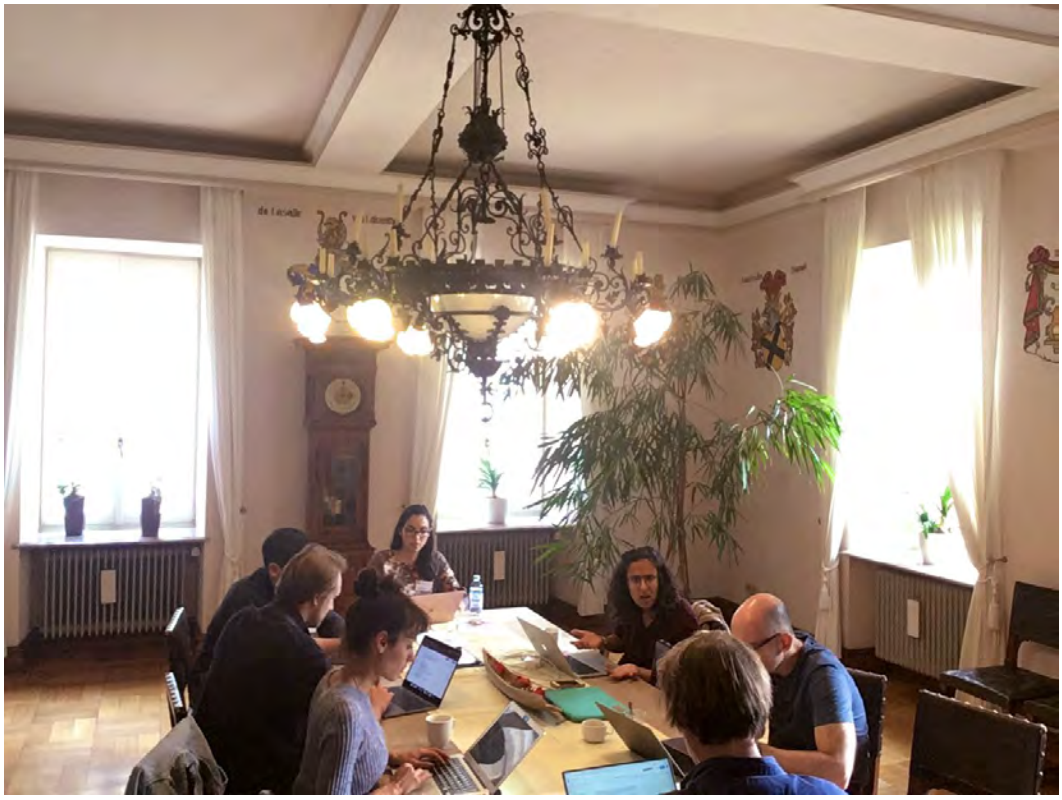

Fig. 6.17

"Working with these amazing people @dagstuhl seminar on Social Agents for Teamwork and Group Interactions" Twitter post by 19411 Dagstuhl Seminar participant Patrícia Alves-Oliveira. https://twitter.com/p_alvesoliveira/status/1182351913632841728. Photo courtesy of Patrícia Alves-Oliveira.

6

6.63 Algorithms and Complexity in Phylogenetics

Organizers: Magnus Bordewich, Britta Dorn, Simone Linz, and Rolf Niedermeier

Seminar No. 19443

Date: October 27–31, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.10.134

Participants: Allan Bai, Magnus Bordewich, Laurent Bulteau, Janosch Döcker, Britta Dorn, Anne-Sophie Himmel, Katharina T. Huber, Mark E. L. Jones, Steven Kelk, Christian Komusiewicz, Simone Linz, Marefatollah Mansouri, Catherine McCartin, Vincent Moulton, André Nichterlein, Rolf Niedermeier, Megan Owen, Charles Semple, Katherine St. John, Leen Stougie, Till Tantau, Nihan Tokaç, Alexandru Tomescu, Leo van Iersel, Mathias Weller, Kristina Wicke, Norbert Zeh

Disentangling the evolutionary relationships between species dates back at least to Charles Darwin and his voyage on board the Beagle. Ever since, the research area of phylogenetics focusses on the reconstruction and analysis of rooted leaf-labeled trees, called phylogenetic (evolutionary) trees, to unravel ancestral relationships between entities like species, languages, and viruses. However, processes such as horizontal gene transfer and hybridization challenge the model of a phylogenetic tree since they result in mosaic patterns of relationships that cannot be represented by a single tree. Indeed, it is now widely acknowledged that rooted leaf-labeled digraphs with underlying cycles, called phylogenetic networks, are better suited to represent evolutionary histories.

Biological questions and applications motivate much of the research in phylogenetics. Nevertheless, most of the software that is routinely used by evolutionary biologists has its roots in theoretical research areas which include algorithms, computational complexity, graph theory, algebra, and probability theory. With a shift from phylogenetic trees towards more complex graphs, the development of new algorithms for phylogenetic networks is currently an active area of research that requires deep insight from computer science and mathematics.

The objective of the seminar was to facilitate interactions between the two research communities of (i) computational and mathematical phylogenetics and (ii) theoretical computer science with a focus on algorithms and complexity. Specifically, its goal was to advance the development of novel algorithms (with provable performance guarantee) to reconstruct and analyze phylogenetic networks that are grounded in techniques from theoretical computer science such as parameterized and approximation algorithms.

This four-day seminar brought together 27 researchers from ten countries, whose research spans theoretical computer science and algorithms, (discrete) mathematics, and computational and mathematical phylogenetics. The seminar program included six overview talks, nine research talks (one of which via Skype), a rump session for short five-minute contributions, and slots for discussions and group work on open problems. More specifically, the overview talks provided introductions to techniques and current trends in parameterized algorithms, combinatorial decompositions, and enumeration algorithms on one hand, and introductions to spaces of phylogenetic trees and networks, and the reconstruction of networks from smaller networks and trees on the other hand. Additionally, each overview talk included open questions and challenges that provided a foundation for discussions and group work throughout the week. The research talks, of which three were given by postgraduate students, covered topical streams of research, including phylogenetic split theory, the placement of phylogenetic problems in higher classes of the polynomial hierarchy, new insight into the popular so-called Tree Containment problem, and phylogenetic diversity and biodiversity indices. Moreover, five working groups were formed on the second day of the seminar. While the research projects that were initiated in these groups are ongoing, some groups obtained first results during the seminar that were presented on the last day.

By building on initially existing synergies between the two research communities, the seminar has taken a leap towards developing new and fostering existing collaborations between both communities. Collaborative work was encouraged and put into practice over formal and informal discussions as well as three group work sessions. Since a significant number of open problems in phylogenetics require the combined expertise of experts in phylogenetics and theoretical computer science, we expect the collaborations formed at Schloss Dagstuhl to make progress on

problems across the traditional discipline boundaries and, ideally, lead to joint peer-reviewed journal or conference publications.

To conclude, this seminar has acknowledged that exchange and connection between the two research communities of theoretical computer science and phylogenetics is fruitful for both sides. Techniques and methods from algorithms and complexity as well as theoretical considerations in general enable, account for, and foster new insights in problems from phylogenetics. Conversely, the specific features and problem structures appearing in the context of phylogenetic trees and networks provide novel theoretical challenges and new directions for foundational research in algorithms and computational complexity.

We thank all participants for their contributions and for openly sharing their ideas and research questions that led to a positive working atmosphere and many discussions throughout the seminar. Furthermore, we sincerely thank the team of Schloss Dagstuhl for their excellent support and communication as well as for providing an enjoyable seminar environment.

Fig. 6.18
"Building topic clusters in the context of Bots in Software Engineering at @dagstuhl #BOTse seminar, which are discussed afterwards in breakout sessions. So many exciting problems to work on!" Twitter post by 19471 Dagstuhl Seminar participant Christoph Matthies. https://twitter.com/chrisma0/status/1196534095460868101. Photo courtesy of Christoph Matthies.

6

6.64 Biggest Failures in Security

Organizers: Frederik Armknecht, Ingrid Verbauwhede, Melanie Volkamer, and Moti Yung

Seminar No. 19451

Date: November 3–8, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.11.1

Participants: Tigist Abera, Magnus Almgren, Frederik Armknecht, Daniel J. Bernstein, Sarani Bhattacharya, Robert Biddle, Begül Bilgin, Dominik Brodowski, Marc C. Dacier, Hervé Debar, Fabio di Franco, Sven Dietrich, Felix Freiling, Daisuke Fujimoto, Lucy Hunt, Ghassan Karame, Stefan Katzenbeisser, Florian Kerschbaum, Youngwoo Kim, Tanja Lange, Gabriele Lenzini, Olivier Levillain, Volkmar Lotz, Michael Meier, Joachim Meyer, Vasily Mikhalev, Christian Müller, Sebastian Pape, Michalis Polychronakis, Kai Rannenberg, Ahmad-Reza Sadeghi, Kazue Sako, Martina Angela Sasse, Stephan Somogyi, Christoph Sorge, Borce Stojkovski, Ingrid Verbauwhede, Melanie Volkamer, Edgar Weippl, Lennert Wouters, Yuval Yarom, Moti Yung

General Introduction

In the present era of ubiquitous digitalization, security is a concern for everyone. Consequently, it evolved as one of the most important fields in computer science. However, one may get the impression that the situation is hopeless. Nearly on a daily basis, reports of new security problems and cyberattacks are published. Thus, one has to admit that despite the huge efforts continuously invested since many decades, securing IT systems remains an open challenge for community and industry.

One of the main reasons is that the variety and complexity of IT systems keeps increasing, making it practically impossible for security experts to grasp the full system. This results into the development of independent and isolated security solutions that at best can close some specific security holes. Summing up, security requires to solve an increasing number of inter- and intradisciplinary challenges while current approaches are not sufficiently effective. The aim of this seminar was to gain an interdisciplinary view on security and to identify new strategies for comprehensively securing IT systems.

Goals

The goals of the seminar was to address the following main challenges and to commonly discuss solution strategies:

Challenge 1: Interdisciplinarity The topic of security is getting more and more complex and already understanding the state-of-the-art within one discipline is highly challenging. This makes it practically impossible to understand the problems and constraints from other disciplines. Moreover, different disciplines often have their own methods and "culture". From our experience, working with colleagues from other disciplines requires at the beginning an enormous effort to understand each other. The complexity grows even further when more than two disciplines are involved.

Challenge 2: Variety of Problems In each discipline, a variety of problems do exist. Naturally, researchers have to single out specific problems that they work on instead of aiming for comprehensive solutions. The selection of problems usually depends on several factors, e.g., background of the researcher, topicality of the subject, etc. Most often, researchers aim for solving very specific problems rather than coming up with more comprehensive solutions. Moreover, the selection is driven by interdisciplinary factors.

For sure, interdisciplinary research does exist already. However, it is mostly restricted to address very few disciplines and has been rather bottom-up by focusing on very specific problems. Instead, the scope of the seminar was to aim for a *broad top-down approach*. To this end, the focus was on the following questions:

- What are the main recurring reasons within disciplines why security solutions fail, i.e., the biggest failures? (Top View)
- How do these failures impact solutions developed in other sub-disciplines? (Broad View)
- What are possible strategies to solve these problems?

Structure

The seminar was structured accordingly. Before the seminar, a survey was conducted where the participants have been asked, what they consider to be biggest failures in security. The list of participants was composed of experts from different, selected sub-fields who were encouraged to explain the main challenges in their field to the audience. Here, ample opportunities for

discussions have been provided. That is, instead of having many different talks back-to-back, we had several overview talks from different fields within the first few days. Afterwards, the whole audience commonly identified three topics to be further investigated in separate working groups:

1. The process and role of certifications
2. The human factor in security
3. The education of the society in security

These subgroups met in parallel and worked on specific questions. The remaining days were composed of workgroup meetings and individual talks. At the end of the seminar, the workgroups reported to the whole audience their findings.

This report summarizes the finding of the survey (Section 3 of the full report), the topics of the individual talks (Section 4 of the full report), and also the findings of the individual workgroups (Section 5 of the full report).

Fig. 6.19
"#Bots at the #Dagstuhl seminar on bots for SE. #BOTse" Twitter post by 19471 Dagstuhl Seminar participant Andreas Schreiber. https://twitter.com/onyame/status/1197514617456603136. Photo courtesy of A. Schreiber/DLR (CC-BY 3.0).

6.65 Machine Learning Meets Visualization to Make Artificial Intelligence Interpretable

Organizers: Enrico Bertini, Peer-Timo Bremer, Daniela Oelke, and Jayaraman Thiagarajan

Seminar No. 19452

Date: November 3–8, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.11.24

Participants: Rushil Anirudh, Enrico Bertini, Alexander Binder, Peer-Timo Bremer, Mennatallah El-Assady, Sorelle Friedler, Beatrice Gobbo, Nikou Guennemann, Nathan Hodas, Daniel A. Keim, Been Kim, Gordon Kindlmann, Sebastian Lapuschkin, Heike Leitte, Yao Ming, Elisabeth Moore, Daniela Oelke, Steve Petruzza, Maria Riveiro, Carlos E. Scheidegger, Sarah Schulz, Hendrik Strobelt, Simone Stumpf, Jayaraman Thiagarajan, Jarke J. van Wijk

The recent advances in machine learning (ML) have led to unprecedented successes in areas such as computer vision and natural language processing. In the future, these technologies promise to revolutionize everything ranging from science and engineering to social studies and policy making. However, one of the fundamental challenges in making these technologies useful, usable, reliable and trustworthy is that they are all driven by extremely complex models for which it is impossible to derive simple (closed-format) descriptions and explanations. Mapping decisions from a learned model to human perceptions and understanding of that world is very challenging. Consequently, a detailed understanding of the behavior of these AI systems remains elusive, thus making it difficult (and sometimes impossible) to distinguish between actual knowledge and artifacts in the data presented to a model. This fundamental limitation should be addressed in order to support model optimization, understand risks, disseminate decisions and findings, and most importantly to promote trust.

While this grand challenge can be partially addressed by designing novel theoretical techniques to validate and reason about models/data, in practice, they are found to be grossly insufficient due to our inability to translate the requirements from real-world applications into tractable mathematical formulations. For example, concerns about AI systems (e.g., biases) are intimately connected to several human factors such as how information is perceived, cognitive biases, etc. This crucial gap has given rise to the field of *interpretable machine learning*, which at its core is concerned with providing a human user better understanding of the model's logic and behavior. In recent years, the machine learning community, as well as virtually all application areas, have seen a rapid expansion of research efforts in interpretability and related topics. In the process, visualization, or more generally interactive systems, have become a key component of these efforts since they provide one avenue to exploit expert intuition and hypothesis-driven exploration. However, due to the unprecedented speed with which the field is currently progressing, it is difficult for the various communities to maintain a cohesive picture of the state of the art and the open challenges; especially given the extreme diversity of the research areas involved.

The focus of this Dagstuhl Seminar was to convene various stakeholders to jointly discuss needs, characterize open research challenges, and propose a joint research agenda. In particular, three different stakeholders were engaged in this seminar: application experts with unmet needs and practical problems; machine learning researchers who are the main source of theoretical advances; and visualization and HCI experts that can devise intuitive representations and exploration frameworks for practical solutions. Through this seminar, the group of researchers discussed the state of practice, identified crucial gaps and research challenges, and formulated a joint research agenda to guide research in interpretable ML.

Program Overview

The main goal of this Dagstuhl seminar was to discuss the current state and future research directions of interpretable Machine Learning. Because two different scientific communities met, the Machine Learning community and the Visualization community, we started the seminar by discussing and defining important terms and concepts of the field. Afterwards, we split up into working groups to collect answers to the following questions: "*Who needs interpretable machine learning? For what task is it needed? Why is it needed?*". This step was then followed by a

series of application lightning talks (please refer to the abstracts below for details).

On the second day, we had two overview talks, one covering the machine learning perspective on interpretability, and the other one the visualization perspective on the topic. Afterwards, we built working groups to collect research challenges from the presented applications and beyond.

The third day was dedicated to clustering the research challenges into priority research directions. The following priority research directions were identified:

- Interpreting Learned Features and Learning Interpretable Features
- Evaluation of Interpretability Methods
- Evaluation and Model Comparison with Interpretable Machine Learning
- Uncertainty
- Visual Encoding and Interactivity
- Interpretability Methods
- Human-Centered Design

On Thursday, the priority research directions were further detailed in working groups. We had two rounds of working groups in which 3, respectively 4, priority research challenges were discussed in parallel by the groups according to the following aspects: problem statement, sub-challenges, example applications, and related priority research directions. Furthermore, all research challenges were mapped into descriptive axes of the problem space and the solution space.

On the last day, we designed an overview diagram that helps to communicate the result to the larger scientific community.

6.66 Conversational Search

Organizers: Avishek Anand, Lawrence Cavedon, Hideo Joho, Mark Sanderson, and Benno Stein

Seminar No. 19461

Date: November 10–15, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.11.34

Participants: Khalid Al-Khatib, Avishek Anand, Elisabeth André, Jaime Arguello, Leif Azzopardi, Krisztian Balog, Nicholas J. Belkin, Robert Capra, Lawrence Cavedon, Leigh Clark, Phil Cohen, Ido Dagan, Arjen P. de Vries, Ondrej Dusek, Jens Edlund, Lucie Flekova, Bernd Fröhlich, Norbert Fuhr, Ujwal Gadiraju, Matthias Hagen, Claudia Hauff, Gerhard Heyer, Hideo Joho, Rosie Jones, Ronald M. Kaplan, Mounia Lalmas, Jurek Leonhardt, David Maxwell, Sharon Oviatt, Martin Potthast, Filip Radlinski, Rishiraj Saha Roy, Mark Sanderson, Ruihua Song, Laure Soulier, Benno Stein, Markus Strohmaier, Idan Szpektor, Jaime Teevan, Johanne Trippas, Svitlana Vakulenko, Henning Wachsmuth, Emine Yilmaz, Hamed Zamani

Background and Motivation

The Conversational Search Paradigm promises to satisfy information needs using human-like dialogs, be it in spoken or in written form. This kind of "information-providing dialogs" will increasingly happen enpassant and spontaneously, probably triggered by smart objects with which we are surrounded such as intelligent assistants such as Amazon Alexa, Apple Siri, Google Assistant, and Microsoft Cortana, domestic appliances, environmental control devices, toys, or autonomous robots and vehicles. The outlined development marks a paradigm shift for information technology, and the key question(s) is (are):

What does Conversational Search mean and how to make the most of it–given the possibilities and the restrictions that come along with this paradigm?

Currently, our understanding is still too limited to exploit the Conversational Search Paradigm for effectively satisfying the existing diversity of information needs. Hence, with this first Dagstuhl Seminar on Conversational Search we intend to bring together leading researchers from relevant communities to understand and to analyze this promising retrieval paradigm and its future from different angles.

Among others, we expect to discuss issues related to interactivity, result presentation, clarification, user models, and evaluation, but also search behavior that can lead into a human-machine debate or an argumentation related to the information need in question.

Moreover, we expect to define, shape, and formalize a set of corresponding problems to be addressed, as well as to highlight associated challenges that are expected to come in the form of multiple modalities and multiple users. Correspondingly, we intend to define a roadmap for establishing a new interdisciplinary research community around Conversational Search, for which the seminar will serve as a prominent scientific event, with hopefully many future events to come.

Seminar Program

A 5-day program of the seminar consisted of six introductory and background sessions, three visionary talk sessions, one industry talk session, and nine breakout discussion and reporting sessions. The seminar also had three social events during the program. The detail program of the seminar is available online. [60]

Pre-Seminar Activities.

Prior to the seminar, participants were asked to provide inputs to the following questions and request:

1. What are your ideas of the "ultimate" conversational search system?
2. Please list, from the perspective of your research field, important open questions or challenges in conversational search.
3. What are the three papers a PhD student in conversational search should read and why?

From the survey, the following topics were initially emerged as interests of participants. Many of these topics were discussed at length in the seminar.

[60] https://www.dagstuhl.de/schedules/19461.pdf

- Understanding nature of information seeking in the context of conversational agents
- Modelling problems in conversational search
- Clarification and explanation
- Evaluation in conversational search systems
- Ethics and privacy in conversational systems
- Extending the problem space beyond the search interface and Q/A

Another outcome of the above pre-seminar questions was a compilation of recommended reading list to gain a solid understanding of topics and technologies that were related to the research on Conversational Search. The reading list is provided in Section 5 of the full report.

Invited Talks.

One of the main goals and challenges of this seminar was to bring a broad range of researchers together to discuss Conversational Search, which required to establish common terminologies among participants. Therefore, we had a series of 18 iinvited talk throughout the seminar program to facilitate the understanding and discussion of conversational search and its potential enabling technologies. The main part of this report includes the abstract of all talks.

Working Groups

In the afternoon of Day 2, initial working groups were formed based on the inputs to the pre-seminar questionnaires, introductory and background talks, and discussions among participants. On Day 3, the grouping was revisited and updated, and, eventually, the following seven groups were formed to focus on topics such as the definition, evaluation, modelling, explanation, scenarios, applications, and prototype of Conversational Search.

- Defining Conversational Search
- Evaluating Conversational Search
- Modeling in Conversational Search
- Argumentation and Explanation
- Scenarios that Invite Conversational Search
- Conversation Search for Learning Technologies
- Common Conversational Community Prototype: Scholarly Conversational Assistant

We have summarized the working groups' outcomes in the following. Please refer to the main part of this report for the full description of the findings.

Defining Conversational Search. This group aimed to bring structure and common terminology to the different aspects of conversational search systems that characterise the field. After reviewing existing concepts such as Conversational Answer Retrieval and Conversational Information Seeking, the group offers a typology of Conversational Search systems via functional extensions of information retrieval systems, chatbots, and dialogue systems. The group further elaborates the attributes of Conversational Search by discussing its dimensions and desirable additional properties. Their report suggests types of systems that should not be confused as conversational search systems.

Evaluating Conversational Search. This group addressed how to determine the quality of conversational search for evaluation. They first describe the complexity of conversation between search systems and users, followed by a discussion of the motivation and broader tasks as the context of conversational search that can inform the design of conversational search evaluation. The group also surveys 12 recent tasks and datasets that can be exploited for evaluation of conversational search. Their report presents several dimensions in the evaluation such as User, Retrieval, and Dialog, and suggests that the dimensions might have an overlap with those of Interactive Information Retrieval.

Modeling Conversational Search. This group addressed what should be modeled from the real world to achieve a successful conversational search and how. They explain why a range of concepts and variables such as capabilities and resources of systems, beliefs and goals of users, history and current status of process, and search topics and tasks should be considered to advance understanding between systems and users in the context of Conversational Search. The group points out that the options the current search engines present to users can be too broad in conversational interaction. They suggest that a deeper modeling of users' beliefs and wants, development of reflective mechanisms, and finding a good balance between macroscopic and microscopic modeling are promising directions for future research.

Argumentation and Explanation. Motivated by inevitable influences made to users due to the course of actions and choices of search engines, this group explored how the research on argumentation and explanation can mitigate some of potential biases generated during conversational search processes, and facilitate users' decision-making by acknowledging different viewpoints of a topic. The group suggests a research scheme that consists of three layers: a conversational layer, a demographics layer, and a topic layer. Also, their report explains that argumentation and explanation should be carefully considered when search systems (1) select, (2) arrange, and (3) phrase the information presented to the users. Creating an annotated corpus with these elements is the next step in this direction.

Scenarios for Conversational Search. This group aimed to identify scenarios that invite conversational search, given that natural language conversation might not always be the best way to search in some context. Their report summarises that modality and task of search are the two cases where conversational search might make sense. Modality can be determined by a situation such as driving or cooking, or devices at hand such as a smartwatch or AR/VR systems. As for the task, the group explains that the usefulness of conversational search increases as the level of exploration and complexity increases in tasks. On the other hand, simple information needs, highly ambiguous situations, or very social situations might not be the bast case for conversational search. Proposed scenarios include a mechanic fixing a machine, two people searching for a place for dinner, learning about a recent medical diagnosis, and following up on a news article to learn more.

Conversation Search for Learning Technologies. This group discussed the implication of conversational search from learning perspectives. The report highlights the importance of search technologies in lifelong learning and education, and the challenges due to complexity of learning processes. The group points out that multimodal interaction is particularly useful for educational and learning goals since it can support students with diverse background. Based on these discussions, the report suggests several research directions including extension of modalities to speech, writing, touch, gaze, and gesturing, integration of multimodal inputs/outputs with existing IR techniques, and application of multimodal signals to user modelling.

Common Conversational Community Prototype: Scholarly Conversational Assistant. This group proposed to develop and operate a prototype conversational search system for scholarly activities as academic resources that support research on conversational search. Example activities include finding articles for a new area of interest, planning sessions to attend in a conference, or determining conference PC members. The proposed prototype is expected to serve as a useful search tool, a means to create datasets, and a platform for community-based evaluation campaigns. The group outlined also a road map of the development of a Scholarly Conversational Assistant. The report includes a set of software platforms, scientific IR tools, open source conversational agents, and data collections that can be exploited in conversational search work.

Conclusions

Leading researchers from diverse domains in academia and industries investigated the essence, attributes, architecture, applications, challenges, and opportunities of Conversational Search in the seminar. One clear signal from the seminar is that research opportunities to advance Conversational Search are available to many areas and collaboration in an interdisciplinary community is essential to achieve the goal. This report should serve as one of the main sources to facilitate such diverse research programs on Conversational Search.

6.67 BOTse: Bots in Software Engineering

Organizers: James D. Herbsleb, Carolyn Penstein Rosé, Alexander Serebrenik, Margaret-Anne Storey, and Thomas Zimmermann

Seminar No. 19471

Date: November 17–22, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.11.84

Participants: Shivali Agarwal, Ireti Amojo, Ivan Beschastnikh, Kelly Blincoe, Nick Bradley, Fabio Calefato, Nathan Cassee, Jacek Czerwonka, Antske Fokkens, Denae Ford, Thomas Fritz, Marco Gerosa, Daniel Graziotin, Sonia Haiduc, James D. Herbsleb, Abram Hindle, Akinori Ihara, Minha Lee, Philipp Leitner, Marin Litoiu, Walid Maalej, Christoph Matthies, Marie-Francine Moens, Martin Monperrus, Claudia Müller-Birn, Nicole Novielli, Ayushi Rastogi, Paige Rodeghero, Carolyn Penstein Rosé, Anita Sarma, Andreas Schreiber, Alexander Serebrenik, Emad Shihab, Arfon Smith, Igor Steinmacher, Margaret-Anne Storey, David R. Traum, Christoph Treude, Bogdan Vasilescu, Stefan Wagner, Mairieli Wessel, Jie Zhang, Thomas Zimmermann

This Dagstuhl seminar brought researchers and practitioners together from multiple research communities with disparate views of what bots are and what they can do for software engineering. The goals were to understand how bots are used today, how they could be used in innovative ways in the future, how the use of bots can be compared and synthesized, and to identify and share risks and challenges that may emerge from using bots in practice.

Bots, often called chatbots, are considered by some to be computer programs that provide a conversational style interface for interacting with software services, while others consider bots to be any semi-autonomous software service that may or may not take on a human-like persona.

Regardless of the definition of what makes a bot a bot, bots are found in many domains such as shopping, entertainment, education, and personal productivity. In software development, bots are rapidly becoming a *de facto* interface for developers and end users to interact with software services in a myriad of ways: e.g., bots are used to fetch or share information, extract and analyze data, detect and monitor events and activities in communication and social media, connect developers with key stakeholders or with other tools, and provide feedback and recommendations on individual and collaborative tasks.

Through this Dagstuhl Seminar, we aimed to gain important insights on how bots may play a role in improving software development productivity and in enhancing collaborative software development. In particular we discussed how bots, with or without a conversational UI, may play a prominent role in software practice. We gathered literature and resources on how bots can have an impact on development processes, software quality, and on end users. The goal was to channel previously siloed communities and through this confluence forge a common vision and plot next steps that might leverage the variety of expertise and push forward both the research and the practices related to bots. The activities were meant to surface the difficult questions and tensions that arise when one looks beyond what at first blush appears to be a superficial distinction, but in fact touches upon core values and driving questions that define the boundaries between fields.

6.68 Composing Model-Based Analysis Tools

Organizers: Francisco Durán, Robert Heinrich, Diego Pérez-Palacín, Carolyn L. Talcott, and Steffen Zschaler

Seminar No. 19481

Date: November 24–29, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.11.97

Participants: Sofia Ananieva, Kyungmin Bae, Christel Baier, Olivier Barais, Simona Bernardi, Erwan Bousse, Benoit Combemale, Juan De Lara, Francisco Durán, Michalis Famelis, Martin Gogolla, Esther Guerra, Robert Heinrich, Mark Hills, Jean-Marc Jézéquel, Kenneth Johnson, Narges Khakpour, Sandro Koch, Raffaela Mirandola, Sam Owre, Diego Pérez-Palacín, Fiona A. C. Polack, Daniel Ratiu, Arend Rensink, Ralf H. Reussner, Elvinia Riccobene, Bernhard Rumpe, Houari Sahraoui, Patrizia Scandurra, Séverine Sentilles, Marjan Sirjani, Carolyn L. Talcott, Catia Trubiani, Hans Vangheluwe, Dániel Varró, Arthur Vetter, Markus Völter, Marc Zeller, Steffen Zschaler

Quality properties like performance and dependability are key for today's systems. Several techniques have been developed to effectively model quality properties, which allow analyzing these systems. However, the very different nature of these properties has led to the use of different techniques and mostly independent tools. In addition, different tools and techniques can be used for modelling quality depending on the size and complexity of the systems and the available details. For example, for modeling dependability techniques like Fault Trees, Markov Chains, and Reliability Block Diagrams are available. Similarly, a range of analysis techniques are available, including simulations, using numerical, analytical or graphical techniques, and analytical methods.

Although it is worth exploring other techniques and methodologies, model-driven engineering (MDE) seems a promising technique to efficiently design and reason about behavior and quality of systems in various domains. Indeed, it has been very successfully applied to improve the efficiency of software development and analysis in various domains.

Moreover, recent innovations, like the Internet of Things, production automation, and cyber-physical systems, combine several domains such as software, electronics and mechanics. Consequently, also the analyses for each of these individual domains need to be combined to predictively analyze the overall behavior and quality. The composition of systems and their analyses is a challenging but unavoidable issue for today's complex systems. Existing MDE approaches to modeling and analysis are not sufficient to compose modular analyses combining domain-specific languages. First attempts towards composable modular models have been developed in recent years, attempting to compose, not only the structure of models and domain specific modeling languages (DSMLs), but also their dynamic aspects (behavior and semantics). These indeed may be good foundations for building composable modular analyses. However, much work remains ahead.

In this Dagstuhl Seminar, we target more flexibility in MDE by discussing how to modularize and compose models and analyses. This provokes questions from the theoretical computer science and formal methods community – for example, on validity, uncertainties, behavior and property protection/preservation/reflection, and termination of analyses. Traditionally, research on these topics is conducted in the formal methods community isolated from the MDE community. A key objective for bringing together representatives from industry and researchers in the formal methods and software engineering communities is to make progress towards establishing the foundations for a common understanding.

6.69 Diversity, Fairness, and Data-Driven Personalization in (News) Recommender System

Organizers: Abraham Bernstein, Claes De Vreese, Natali Helberger, Wolfgang Schulz, and Katharina A. Zweig

Seminar No. 19482

Date: November 24–29, 2019 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.9.11.117

Participants: Christian Baden, Michael Beam, Abraham Bernstein, Claes De Vreese, Marc Hauer, Lucien Heitz, Natali Helberger, Pascal Jürgens, Christian Katzenbach, Benjamin Kille, Beate Klimkiewicz, Wiebke Loosen, Judith Möller, Goran Radanovic, Wolfgang Schulz, Guy Shani, Nava Tintarev, Suzanne Tolmeijer, Wouter van Atteveldt, Sanne Vrijenhoek, Theresa Züger, Katharina A. Zweig

The Dagstuhl Perspectives Workshop 19482 on Diversity, Fairness, and Data-Driven Personalization in (News) Recommender Systems,[61] took place from November 24 to November 29 at Schloss Dagstuhl in Germany. The goal of the workshop was to bring together researchers from the various disciplines relevant to news recommender systems (computer, communications, legal, and political science) to (1) develop a joint understanding of the issues arising for society with regards to the diversity and fairness of recommender systems, (2) identify the gaps in science, practice and regulation with regards to these topics, and (3) to compile a set of recommendations–in the form of a manifesto–that outlines needed steps from all actors involved to address the societal issues at hand.

Workshop Schedule

The workshop was organized in the following phases:

Welcome and introductions This first phase introduced the workshop goal to the participants and then offered each of them five minutes to introduce their research activities, expertise, their interest in the topic, and research directions they see as relevant to the workshop's topic.

Impulse presentations Given the diversity of the backgrounds of the participants, eight brief stage setting presentations where given. The goal of these was to establish a common ground in terms of relevant questions and common vocabulary.

Topical breakout group discussions Based on the introducing presentations and impulse presentations, the next phase of the workshop was organized around topical breakout groups. Topics discussed included relating fairness to diversity, user desiderata and characteristics, wider societal implications, governance, data requirements, and clustering of research gaps.

Writing sessions The next phase was focused on jointly drafting the manifesto that incorporated recommendations developed from discussions so far and compiling them into a coherent document.

The remainder of this text provides the abstracts of the impulse presentations. The insights resulting from our discussions can be found in the manifesto document, which will be published in due course.

[61] See workshop home page at https://www.dagstuhl.de/19482

6

6.70 Big Graph Processing Systems

Organizers: Angela Bonifati, Alexandru Iosup, Sherif Sakr, and Hannes Voigt

Seminar No. 19491

Date: December 1–6, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.12.1

Participants: Khaled Ammar, Renzo Angles, Walid Aref, Marcelo Arenas, Maciej Besta, Peter A. Boncz, Angela Bonifati, Khuzaima Daudjee, Emanuele Della Valle, Stefania Dumbrava, Olaf Hartig, Bernhard Haslhofer, Tim Hegeman, Jan Hidders, Katja Hose, Adriana Iamnitchi, Alexandru Iosup, Vasiliki Kalavri, Hugo Kapp, Wim Martens, M. Tamer Özsu, Eric Peukert, Stefan Plantikow, Mohamed Ragab, Matei R. Ripeanu, Sherif Sakr, Semih Salihoglu, Christian Schulz, Petra Selmer, Juan F. Sequeda, Joshua Shinavier, Gábor Szárnyas, Riccardo Tommasini, Antonino Tumeo, Alexandru Uta, Ana Lucia Varbanescu, Hannes Voigt, Hsiang-Yun Wu, Nikolay Yakovets, Da Yan, Eiko Yoneki

In memoriam: This seminar is dedicated to the memory of our co-organizer and friend Sherif Sakr (1979-2020), whose unexpected early departure happened a few months after the seminar. Sherif was a leading scientist in the field of Big Data Technologies. We are grateful to him for the time spent together and the joint work preceding and following the seminar. He will be deeply missed.

The world has become more interconnected than ever. Through an advancing wave of technologies and applications, our society is producing and consuming data at an unprecedented scale and complexity. To model the data, graphs offer a general model and mathematical abstraction, in the simplest form based on arbitrary objects (vertices) connected by relationships (edges), with possibly additional information (properties[62]). Graphs enable already a remarkable range of application domains[63], from industry to science, from society to governance, from education to gaming, but their true potential is just beginning to be unlocked. However, the tremendous increase in the size, complexity, and diversity of the graph-structured data and their applications, and the increasing community using graphs to understand and automate the world around us, raises new challenges for computer science. Under these new circumstances, the potential benefits of graph processing could be canceled by the difficulty to understand, create, develop, and automate graph processing for the masses. Focusing on the interplay between graph data, abstractions, systems, performance engineering, and software engineering, this seminar brings together researchers, developers, and practitioners actively working on this topic, to discuss *timely and relevant* open challenges with a main focus on the following topics: trade-off of design decisions of big graph processing systems, high-level graph programming abstractions and graph query languages, the specific requirements for different application domains for benchmarking and graph engineering purposes, systems and ecosystems for graph processing, the fundamental processes and methods leading to the science, design, and engineering of graph processing.

The seminar focused on the following key topics related to big graph processing systems:

Topic 1. Design Decisions of Big Graph Processing Ecosystems: In modern setups, graph-processing is not a self-sustained, independent activity, but rather part of a larger big-data processing ecosystem. Typical examples include the Giraph's deployment in the Facebook MapReduce ecosystem[64], Powergraph[65] in the GraphLab[66] machine learning and data-mining ecosystem, and GraphX[67] in the Apache Spark ecosystem. In general, more alternatives usually mean harder decisions for choice. In practice, with the wide spectrum of big graph processing systems, with different design decisions, that are currently available, it becomes very challenging to decide by intuition which system is the most adequate for a given application requirements, workload, or the underlying ecosystem. Making such decisions requires significant knowledge about the graph

[62] M. Junghanns et al., "Analyzing Extended Property Graphs with Apache Flink," NDA'16
[63] L. da Fontoura Costa et al. (2008) Analyzing and Modeling Real-World Phenomena with Complex Networks: A Survey of Applications, ArXiv Physics and Society, 2008. https://arxiv.org/abs/0711.3199v3 This study identifies tens of application domains for graph processing.
[64] Ching et al., One Trillion Edges: Graph Processing at Facebook-Scale, VLDB '15.
[65] Gonzalez et al., PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs, OSDI '12.
[66] Low et al., Distributed GraphLab: a framework for machine learning and data mining in the cloud, VLDB '12.
[67] Gonzalez et al., GraphX: Graph Processing in a Distributed Dataflow Framework., OSDI '12.

complexity, graph size, world requirements, and even the implementation details of the various available systems. Currently, we lack the fundamental models to understand and quantitatively analyze, estimate, and describe the complexity of big graph processing jobs. In addition, there is no understanding on the relationship between the graph complexity and the computational complexity of big graph processing jobs. Therefore, we need a clear understanding for the impact and the trade-offs of the various decisions (e.g., centralized vs distributed, partitioning strategy, programming model, graph representation model, memory storage vs disk storage) in order to effectively guide the developers of big graph processing applications.

Topic 2. High-Level Graph Processing Abstractions: While imperative programming models, such as *vertex-centric* or *edge-centric* programming models, are popular, they are lacking a high-level exposition to the end user. This way the end user is required more technical programming, which limits the end user productivity in building graph processing pipelines. In contrast, graph query languages build on more high-level, declarative constructs. Query language abstraction give more power to the less technical user and allow for extensive performance optimization by the underlying graph processing system. Current graph query languages, however, lack the power required in many graph analytics use cases. To increase the power of graph processing systems and foster the usage of graph analytics in applications, we need to design high-level graph processing abstractions. It is currently completely open how future declarative graph processing abstractions could look like, which the best level of abstraction is, how abstraction for analytics integrate with existing graph query languages, and we can evaluate new graph processing abstractions regarding utility, simplicity, expressiveness, and optimization potential.

Topic 3. Performance and Scalability Evaluation: Traditionally, performance and scalability are measures of efficiency, contrasting the ability of systems to utilize resources: FLOPS, throughput (e.g., EVPS), or speedup (i.e., compared to either a single-node, or a sequential implementation). Such metrics are difficult to apply for graph processing, especially since performance is non-trivially dependent on *platform*, *algorithm*, and *dataset* (i.e., the PAD triangle[68]). Therefore, many important questions arise: *how to compare the performance of graph-processing systems?, how to define scalability?, should one compare largely different systems, e.g., a distributed, heterogeneous system with a highly-tuned, hand-written sequential implementation?, how to design a framework for reproducible performance evaluation?*. Moreover, running graph-processing workloads in the cloud leverages additional challenges. First, we would like to understand whether the intrinsic cloud elasticity could be harnessed for graph processing. Second, clouds are known to be impacted by large degrees of performance variability due to colocation and virtualization overheads. Studying the impact of cloud performance variability onto graph-processing workloads is another topic of interest. Such performance-related issues are key to identify, design, and build upon widely recognized benchmarks for graph processing.

For each topic, the discussion also considered specific and general applications of graph processing, at various volume, velocity, and other dimensions.

The seminar brought together over 40 diverse and high quality researchers with core expertise from two generally distinct communities, data management and (large-scale) computer systems. The seminar was successful, and addressed in particular topics around graph processing systems: ecosystems, abstractions and other fundamental theory, and performance. To this end, we structured the seminar as follows:

1. Prior to the seminar, the co-organizers have contacted each participant, eliciting commitment for one or several topics, and ideas for key elements of the discussion.
2. During the first day of the seminar, the morning was dedicated to short presentations by each participant, and a long break-out session per topic. The former allowed the participants to better understand each other's core ideas and keywords, to identify synergies and to find experts for keywords not entirely familiar.
3. For the next two days, each morning challenged at least one half the participants with a tutorial given by a leading expert from the *other* community, then proceeded with break-out sessions organized per topic, and ended with a plenary session to share the main ideas. The tutorials were given by Tamer Özsu on "Graph Processing: A Panoramic View and Some Open Problems", on behalf of the data management community, and by Antonino Tumeo on "Big Graph Processing: The System Perspective", on behalf of the systems community. The main results of these two days of intense work were making terminology more uniform across the participants, and the core ideas about challenges (open problems), directions for long-term research, and identification of concrete short-term plans for continuation.
4. During the last day of the seminar, the participants finalized the immediate conclusions of the seminar (see Section "In Conclusion: Challenges and Future Directions for Big Graph Processing Systems"), and agreed on the plans for continuation.

[68] Guo et al., How well do graph-processing platforms perform? an empirical performance evaluation and analysis, IPDPS '14.

6

6.71 Future Automotive HW/SW Platform Design

Organizers: Dirk Ziegenbein, Selma Saidi, Xiaobo Sharon Hu, and Sebastian Steinhorst

Seminar No. 19502

Date: December 8–11, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.12.28

Participants: Bart Besselink, Alessandro Biondi, Jerónimo Castrillón-Mazo, Lulu Chan, Wanli Chang, Thidapat Chantem, Jyotirmoy Deshmukh, Rolf Ernst, Sabine Glesner, Masaki Gondo, Baik Hoh, Oliver Kopp, Peter Gorm Larsen, Mark Lawford, Roland Leißa, Chung-Wei Lin, Martina Maggio, Albrecht Mayer, Frank Mueller, Philipp Mundhenk, Alessandra Nardi, Moritz Neukirchner, Philipp Obergfell, Maximilian Odendahl, Zhu Qi, Eduardo Quinones, Sophie Quinton, Selma Saidi, Ignacio Sañudo Olmedo, Lea Schönberger, Lukas Sommer, Wilfried Steiner, Seyhan Uçar, Dirk Ziegenbein

Driven by new functionality and applications (such as automated driving and vehicle-to-X-connectivity) and fueled by the entry of new players from the IT industry, automotive systems are currently undergoing a radical shift in the way they are designed, implemented, and deployed. The trend towards automation and connectivity imposes an increased complexity and requires unprecedented computing resources, while, at the same time, the demanding requirements regarding cost-efficiency and dependability still need to be fulfilled. One of the most visible changes is the integration of formerly separated function domains onto centralized computing platforms. This leads to a heterogeneous mix of applications with different models of computation (e.g., control, stream processing, and cognition) on heterogeneous, specialized hardware platforms (comprising, e.g., application cores, safety cores, GPUs, deep learning accelerators) to accommodate advanced functionalities such as automated driving and on-line optimization of operating strategies for electrified powertrains.

The adoption of these novel heterogeneous platforms raises several challenges. In particular, many of their components stem from embedded consumer devices and have never been designed for application in safety-critical real-time systems. Therefore, while their computational capabilities are well understood, there is an increased need to comprehend these platforms from the perspective of extra-functional requirements such as predictability, determinism, and freedom-from-interference. This process deeply impacts the core design aspects of automotive E/E architectures and heavily challenges established methods and methodologies in HW/SW automotive design.

The goal of this Dagstuhl Seminar was to gather researchers and practitioners from academia and industry to discuss key industrial challenges, existing solutions and research directions in the HW/SW design of future automotive platforms. The seminar focussed, in particular, on

- predictability of systems regarding extra-functional properties,
- safe integration of hardware and software components and
- programmability and optimization of emerging heterogeneous platforms.

These inter-dependent challenges require the interaction between multiple disciplines, combining resource-constrained embedded, cyber-physical, and real-time aspects. Another important aspect of the seminar was to provide insight into novel automotive functionalities (such as automated driving, online optimization, or over-the-air-update) and their software architectures and requirements as well as into the HW/SW platforms they are executed on.

The seminar provided a unique opportunity for participants from the automotive industry to present their challenges and constraints and receive feedback and ideas from academia. At the same time, it allowed researchers to confront their own ideas and/or solutions with industrial reality and together identify new research directions in order to make an impact in the automotive industry.

Organization of the seminar

The seminar took place from 8th to 11th December 2019. The seminar started with an overview of current trends and challenges in the design of future automotive HW/SW platforms by the organizers. After that the agenda was structured along the previously mentioned challenges. Monday's talk sessions were focused on dependability and predictability of HW/SW systems. The sessions on Tuesday dealt with the safe integration of heterogeneous software applications covering aspects of software architectures, networks and cyber-physical systems in the automotive domain and touched societal issues as well.

On Wednesday, the talks focused on the programmability and optimization of heterogeneous platforms. All talks were restricted to 15 minutes, leaving ample time for discussions as well as breakout sessions on the following topics:

- Modeling hardware and software dependencies
- Weakly hard real-time models
- Machine learning in cyber-physical systems
- HW/SW architecture exchange
- Benchmarking efforts for future HW/SW platforms
- Modularizing control systems
- Automotive software lifecycle
- Programming vs. execution models

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

Outcome

The seminar succeeded in bringing together participants from different communities who were engaged in very intensive, interdisciplinary group discussions. Not surprisingly, many participants stated that they were able to learn a lot from adjacent fields. As many of the industrial challenges at hand require interdisplinary approaches, the organizers consider this a significant success of the seminar. One example that became evident during the course of the seminar was that terms like execution model are quite differently used in e.g. the high performance computing domain and in the embedded systems community. A group formed in one of the breakout sessions intends to write a whitepaper on unifying terminology and formulating a common understanding of the different layers of models used in designing automotive HW/SW systems. A first follow-up meeting already took place in February 2020.

Several industrial presentations gave valuable insights in the industrial state-of-the-practice and outlined challenges for future research. A very good example for this was the breakout session “HW/SW Architecture Exchange” which discussed current architectural patterns and open challenges in the context of designing dependable systems and achieving deterministic behavior on heterogeneous high-performance HW platforms.

Another breakout session provided an overview of current automotive benchmarks and performance models that can be used as a basis for research activities. This session also raised the awareness that industry needs to be more active in providing relevant benchmarks in order to enable researchers to validate the industrial viability of their solutions.

Overall, the feedback of the participants showed that they made a lot of new contacts in academia and industry and a follow-up seminar in about two years was requested by many participants. The seminar inspired several new collaborations including contributions to the Autonomous Systems Design workshop at DATE 2020, ideas for special sessions at DAC 2020 and ESWEEK 2020 and also a student project on automotive HW/SW platform simulation between a students’ project group and an industrial partner.

6.72 Artificial and Computational Intelligence in Games: Revolutions in Computational Game AI

Organizers: Jialin Liu, Tom Schaul, Pieter Spronck, and Julian Togelius

Seminar No. 19511

Date: December 15–20, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.12.67

Participants: Dan Ashlock, Yngvi Björnsson, Alan Blair, Bruno Bouzy, Cameron Browne, Michael Buro, Duygu Cakmak, Tristan Cazanave, Alex J. Champandard, Michael Cook, Peter I. Cowling, Jakob Foerster, Raluca D. Gaina, Katja Hofmann, Jialin Liu, Setareh Maghsudi, Nantas Nardelli, Mark J. Nelson, Diego Perez Liebana, Mike Preuß, Sebastian Risi, Günter Rudolph, Tómas Philip Rúnarsson, Christoph Salge, Spyridon Samothrakis, Tom Schaul, Jacob Schrum, Emily Short, Chiara F. Sironi, Pieter Spronck, Nathan Sturtevant, Olivier Teytaud, Tommy Thompson, Julian Togelius, Vanessa Volz, Hui Wang, Mark Winands, Georgios N. Yannakakis

The past decade has seen a rapid advent of new technologies in computational game playing. For a long time, artificial intelligence (AI) for 2-player deterministic board games was mainly implemented using tree search, employing the minimax algorithm with alpha-beta pruning, enhanced with some improvements which were often aimed at particular games. This approach worked well for most traditional games, but some games proved to be notoriously hard to tackle in this way. The textbook example of games for which regular tree search is inadequate, is Go.

Ten years ago, the novel technique of Monte Carlo Tree Search (MCTS) became popular, as it was shown that using MCTS, the quality of AI for Go improved significantly, albeit not yet to the level of top-level human players. Many experts predicted that around 2030 Go AI would surpass human-level play. Much to the surprise of many, however, already in 2016 Google's AlphaGo defeated the human world champion in Go, using a combination of MCTS and deep convolutional networks to evaluate Go board positions and perform move selection. The networks were trained using millions of examples of human play, combined with self-play. A short while later, it was demonstrated with AlphaZero that self-play by itself suffices to train the networks to reach the necessary quality.

There is a long history of research into computational game AI for 2-player deterministic board games. However, since the turn of the century computational techniques have also been applied to games of a different nature, such as games for 3 or more players, games with imperfect information, and video games. Such games bring their own challenges, and often need very different approaches for creating game AI. Nevertheless, computational techniques may be applicable. Recent successes have been achieved in the playing of Poker (multiple players, imperfect information) and DotA (team-based video game). Deep learning techniques have been used to teach a game AI to play old Atari video games, and the highly complex game Doom, by only observing the screen.

These computational approaches to AI game playing have been highly successful, and have caused great enthusiasm in researchers and laymen alike. However, while they have opened up new possibilities for implementing strong game AI, they are definitely not the one-size-fits-all solution for all problems in computational game playing. The aim of the seminar was to build upon the foundations laid by the state-of-the-art in computational game playing, and (1) identify for which game AI problems the current state-of-the-art is inadequate or unsuitable, including the reasons why; (2) propose and investigate which improvements to the state-of-the-art may open up ways to apply it to a wider range of game AI problems; and (3) form ideas on which novel techniques may be employed to solve problems for which the current state-of-the-art is simply not suitable.

For the purpose of the seminar, a "game" is considered any simulated environment in which decisions can be taken in order to achieve a particular goal. This includes board games, card games, video games, simulations, and VR/AR applications. Decisions in a game are taken by "players." In multi-player games, the goal is usually to "win" from other players, by reaching a pre-defined victory condition before any other player manages to do so. "Game AI" is a computer-controlled player. Good game AI takes decisions which are highly effective in achieving the goal. Cooperative games are also of interest, where the aim is for the players to work together to share in a victory. We wish to point out that games are often a reflection of some aspects of the real world, and allow investigating those aspects in a risk-free environment – good solutions for problems found in games may therefore have immediate applications in the real world.

The following is a (non-exhaustive) list of challenges to game AI which are hard to deal with using the current state-of-the-art. These challenges formed the basis for the discussions and investigations of the seminar.

- **Determining the limitations of MCTS and deep learning for computational game playing**: The state-of-the-art in computational game playing encompasses Monte-Carlo Tree Search (MCTS) and deep convolutional networks to store

game information. The recent successes of this approach in Go have made MCTS and deep learning the "go-to" techniques for implementing game AI. However, these techniques have many inherent limitations. MCTS needs extraordinary amounts of computer power and is therefore very expensive to use. While it can be parallelized easily, just adding more computer power has diminishing pay-offs. Moreover, there are many games for which MCTS clearly is not a suitable approach, for instance, games with a large branching factor where it is hard to come up with heuristics which pinpoint the branches which are most likely to contain the strong moves. As for deep learning, now that the early enthusiasm has waned a little, the first criticisms of it, which explain its many limitations, are already being published. Gaining insight into the limitations of MCTS and deep learning for game AI implementation will allow us to distinguish those games for which these techniques may be employed for strong game playing from those games for which different approaches are needed.

- **Defining more appropriate game complexity measures**: Game complexity is a measure which is supposed to indicate how difficult it is to implement game AI. It is usually expressed as the number of possible game states in base log_{10}. Beyond a complexity of 100 (10^{100} game states), it is highly unlikely that a game will ever be "solved," i.e., will never be played perfectly. Researchers therefore aim for superhuman rather than perfect play. For a long time Go was considered the pinnacle of complexity in game playing, boasting a game complexity of 360. However, in the game AI domain, games have been researched with a much higher game complexity than Go. Typical examples of such games are:
  - *Arimaa*, a 2-player deterministic, perfect-information board game with a game complexity of 402.
  - *Stratego*, a 2-player deterministic, imperfect-information board game with a game complexity of 535.
  - *StarCraft*, a typical Real-Time-Strategy video game, with a varying game-tree complexity (depending on the parameters of the scenario) which is measured in the tens of thousands.

  The increased complexity of these games stems from multiple factors, such as an increased move complexity (e.g., in Arimaa players always make four moves in sequence), the introduction of imperfect information (e.g., in Stratego at the start of the game the players only know the location of their own pieces), or simply an explosion of pieces, moves, and non-deterministic influences (e.g., most video games). A common belief is that an increase in game complexity also entails an increase in difficulty of creating a game AI; however, previous investigations have shown that high game complexity does not necessarily equate high difficulty for achieving superhuman play. This indicates that "game complexity" might not be the most appropriate complexity measure for games. A theoretical investigation of game features may result in alternative ways to express game complexity, which may better relate to the difficulty of playing the game for an AI. Moreover, a better understanding of what makes a game difficult for an AI, might lead to new insights into how strong game AI can be built.
- **Learning game playing under adverse conditions**: In recent years, most research into game AI has moved towards "learning to play" rather than "implementing an algorithm." Game AI can learn from observing examples of human play (provided a large enough dataset is available) or from self-play. Such learning has lead to strong results in some cases, but in many cases fails under adverse conditions. Examples of such conditions are:
  - Imperfect information, i.e., the results of decisions of the AI depending partly on unknown data.
  - Continuous action spaces, i.e., the AI in principle being allowed to take an unlimited number of decisions in a small time period; thus, an AI not only has to decide what actions it wants to take, but also how many and with which intervals.
  - Deceptive rewards, i.e., situations in which positive results achieved by the AI in the short term, in practice drive it away from the ultimate goal of the game in the long term.

  To implement learning AI for wider classes of games, approaches must be devised to deal with such adverse conditions in systematic ways.
- **Implementing computational game AI for games with 3 or more players**: Most research into game AI is concerned with zero-sum 2-player games. The reason is obvious: in 2-player games, an objective "best move" always exists. With games that involve more than two players, which oppose each other, there often is no obvious "best move." For instance, if there are three players in the game, if two of those players band together, in general the third one will have a very hard time winning the game, even when taking into account that the other two are collaborating. The main problem is that when one player is obviously doing better than the other two, it is to the advantage of the other two to collaborate against the envisioned winner. Therefore, in games with three or more players, it may be advantageous not to play better than the other players, in order not to become the target of a collaborative assault of the opponents. A pessimistic perspective, where the AI assumes that the opponents will actually form a block, will in general lead to much worse play than a perspective wherein the AI tries to form collaborations itself. This means that the AI must incorporate in its reasoning the attitudes of the other players, for instance in the form of player models.

  The topic of AI for games of three or more players has been studied very little – a notable exception being the study of Poker, which is a relatively easy game in this respect considering the simplicity of the required player models and the very small state-space and game-tree complexities. Hundreds of thousands of games for three or more players exist, which by itself means that this challenge needs investigation. Moreover, when translating research findings to real-world challenges, the existence of more than two players is a given in most realistic situations.
- **Implementing AI for games with open-ended action spaces**: Certain classes of games have so-called "open-ended action spaces," i.e., the number of possible actions is basically unlimited. One example of such a type of game is found in interactive fiction: these are puzzle games which the player controls by typing sentences in plain English. While each game only understands a limited set of verbs and nouns, the player is generally unaware of this list. Designers of such games aim to allow the player to give any English command which is reasonable in the circumstances described by the game. Another example of such a type of game is a tabletop role-playing game, in which players are allowed to perform any action at all, and a game master determines (within boundaries of a complex ruleset) what the result of the action is. Creating an AI for either a game master or a player of such a game requires the AI to have at least a basic understanding of the game world to be successful. In practice, studies into

game AI focus almost exclusively on games where the action spaces are closed, which makes regular learning algorithms applicable. For games with open-ended action spaces, as of yet no successful approaches have been found.

- **General Game Playing**: In the domain of General Game Playing (GGP), games are defined by a set of rules, specified in a General Game Description Language (GGDL). The goal of researchers in this domain is to create an artificial intelligence which is able to play such a game, based on only the rules. Yearly competitions are held where researchers pose their AIs against each other in games which are unknown to them at the start of the competition. The state-of-the-art in such competitions is using MCTS, enhanced according to some general assumptions on the types of games that need to be played. This approach is unsurprising, as MCTS does not require knowledge of the game in question in order to do reasonably well. In fact, this approach is so strong and so easy to implement that all competitors use it. The danger of the success of MCTS for GGP is that the research in this area gets stuck at a dead end – the same happened with the research into traditional game AI when for decades researchers only worked on small improvements to minimax and alpha-beta pruning, until MCTS came around to shake things up. It is highly unlikely that a blind and ostensibly "stupid" approach such as MCTS is the end-all of GGP AI implementations. It is therefore of particular interest to investigate novel approaches to GGP, which are not MCTS-based.
- **General Video Game Playing**: General Video Game Playing (GVGP) aims at designing an AI agent which is capable of successfully playing previously-unknown video games without human intervention. In the General Video Game AI (GVGAI) framework, video games are defined by a set of rules, sprites and levels, specified in a Video Game Description Language (VGDL). The VGDL was initially proposed and designed at the 2012 Dagstuhl Seminar on Artificial and Computational Intelligence in Games. The GVGAI framework has been expanded to five different competition tracks: (1) single-player planning, (2) two-player planning, (3) learning (in which no forward model is given), (4) level generation and (5) rule generation. In the planning tracks a forward model of every game is available; MCTS has been the state-of-the-art algorithm in these tracks. However, MCTS is not applicable to the learning track as no forward model is given and thus no simulation of game playing is possible. Deep reinforcement learning is a potential approach for the GVGAI learning track, but has not been investigated yet. Other methods might have potential too. Determining the applicability of different methods to the creation of GVGAI is a novel and topical challenge. Of particular interest in this respect is the creation of an AI for the domain of general Real-Time-Strategy (RTS) games.
- **Computation for human-like play**: Virtually all research into computational game AI focuses on building a game-playing AI which is as strong as possible. Strength can objectively be measured by pitting different AIs against each other. In video-game AI research, it has been recognized that playing strength is, in general, not a major goal – instead, much research in video game AI is aimed at making the AI play in an entertaining, interesting, or human-like manner. MCTS is notoriously unsuitable for selecting moves that are human-like, as it is simply based on finding the best outcome for the game as a whole. However, in situations where humans play against an AI, whether it is for entertainment or training, it is desirable that not only the best moves can be played by the AI, but also those moves which are interesting to explore or are on par with how a human might play. Almost no research has yet been done into computational human-like AI, which makes it a worthy challenge to take on.

The Dagstuhl seminar brought together computer scientists and industry experts with the common goals of gaining a deeper understanding of computational game AI, in particular to determine the limitations to the state of the art, to find new uses for the state-of-the-art, to explore new problems in the domain of computational game AI, and to investigate novel approaches to implementing computational game AI. Industry experts came not only from companies which specifically work in game AI research, but also from companies which use game AI in their products.

During the seminar we not only had discussions which investigate the topics theoretically, but also spent part of the seminar on trying to achieve practical results. We did the same in the 2015 and 2017 seminars, which was met with great enthusiasm and led to some strong follow-ups. As in the previous seminars, these practical sessions were partly to test out new ideas, and partly competition-based, where different approaches were used to implement AI for new problems, which were then compared to each other by running a competition.

What was new for this particular seminar, is that we held expert talks during some of the evenings. These started with one or two experts giving a longer talk (between half an hour and an hour-and-a-half) on one of their specialisms, followed by a longer Q&A session and a discussion. One evening was spent this way on using modern communication media to inform people about research (Tommy Thompson), one evening was spent on the details of DeepMind's AlphaStar (Tom Schaul), and one evening was spent on advanced search techniques in board games (Olivier Teytaud and Tristan Cazanave). These evenings were greatly appreciated by participants, and should remain part of this series of seminars.

Reports on the working groups are presented on the following pages. Many of these working groups have lead to collaborations, which will lead to papers to be published at conferences and in journals in the coming year. All in all, the general impression was that the participants and the organizers found the seminar a great success, and an inspiration for future research.

6.73 Interactive Design and Simulation

Organizers: Thomas A. Grandine, Jörg Peters, and Ulrich Reif

Seminar No. 19512

Date: December 15–20, 2019 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.9.12.115

Participants: Pierre Alliez, Mario Botsch, Fehmi Cirak, Carlotta Giannelli, Thomas A. Grandine, Cindy Marie Grimm, Klaus Hildebrandt, Alec Jacobson, Bert Jüttler, Ladislav Kavan, Adarsh Krishnamurthy, Angela Kunoth, David I. W. Levin, Angelos Mantzaflaris, Dominik Mokriš, Jörg Peters, Francesca Pitolli, Ulrich Reif, Maria Lucia Sampoli, Thomas Takacs, Etienne Vouga, Chris Wojtan, Urška Zore

Dagstuhl Seminar 19512 presented and debated a rich set of techniques for improving algorithms and interfaces for interactive physical simulation, based on geometric and physical models and governed by partial differential equations. The techniques originate in geometry processing, computational geometry, geometric design, and the use of splines in meshing-less and iso-geometry approaches.

Thanks to its diverse roster of participants, with expertise spanning computer science, applied mathematics and engineering, the seminar enabled rare new interactions between academia, industrial and government-sponsored labs and fostered new insights apart from technical considerations. For example, one of the ad hoc discussions centered around the mechanisms and person-to-person considerations that enable transfer of new techniques from academia to industry. Another discussion focused on bridging the divide between geometric modeling and engineering analysis. A third focused on the usage (or lack thereof) of academic open-source libraries. And a fourth elucidated the different error measures that allow or prevent model reduction techniques for non-linear models (e.g. of elasticity) for given applications ranging from animation to product design.

The seminar was well-paced, avoiding densely-packed presentations. An emphasis was placed on time to formulate both specific and long range challenges. The benchmark problem in Section 4 of the full report is an example of specific problems that clarify and contrast the competing approaches and objectives and advertised the different strengths and the synergy of the areas: responses that permit two-sided error bounds, responses based on mathematical reformulation, applying advanced computational geometry and new software packages that leverage hierarchical spline software.

7 Öffentlichkeitsarbeit
Public Relations and Outreach

7.1 Pressemitteilungen und Medienarbeit

7.1 Press Releases and Media Work

Die regelmäßige Erstellung und Herausgabe von Pressemitteilungen dient der verständlichen Verbreitung von aktuellen Informatikthemen. Die Vermittlung des Konzepts von Schloss Dagstuhl ist dabei ebenfalls ein Thema. Pressemitteilungen und Berichterstattungen in diversen Medien – soweit bekannt – sind über das Internetportal von Schloss Dagstuhl[69] 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. 2 150 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. 1 100 Abonnenten verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ gepflegt (derzeit über 620 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[69].

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 2,150 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 1,100 followers. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with more than 620 members), which supports the networking of participants in Dagstuhl Seminars. Additionally, interesting news about Schloss Dagstuhl are announced there.

7.2 Fortbildung

7.2 Educational Training

Lehrerfortbildung

Seit nunmehr fast 30 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.

Die Lehrerfortbildung dauert drei Tage; an jedem Tag werden in jeweils 3-stündigen Vorträgen zwei Informatikthemen vorgestellt. Die intensive Fortbildung richtet sich

Teacher training

Since almost 30 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.

The 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

[69] 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 2019 gibt es auf der Webseite der Veranstaltung[70].

participate. Details on the workshop in 2019 are available at the event webpage[70].

7.3 „Dagstuhler Gespräche“ / *“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 die Veranstaltungsreihe *Dagstuhler Gespräche* weitergeführt. Der interessierten Öffentlichkeit werden hier Themen aus dem breiten Spektrum der Informatik sowie ihre praktische Anwendung im Alltag oder in wirtschaftlichen Prozessen anschaulich in Form eines Impulsvortrages näher gebracht, um danach in einen gemeinsamen Dialog einzusteigen. An den Dagstuhler Gesprächen nehmen Entscheider und Gestalter aus Wirtschaft, Politik und der Informatik teil, aber auch Interessierte aus der Bevölkerung sind herzlich eingeladen.

Für den 17. Mai 2019 konnte der Präsident der Gesellschaft für Informatik e.V., Prof. Dr. Hannes Federrath, als Vortragender gewonnen werden. Unter dem Titel „Was weiß das Internet über mich?“ erklärte er, welche Spuren man im Internet hinterlässt, und ging der Frage nach wie sich die Nutzer im Internet vor Profilbildung schützen können und welche Datenschutzrechte sie im Internet haben.

Noch eine weitere Veranstaltung der Reihe wurde durchgeführt: Am 24. November 2019 trug Prof. Dr. Katharina Zweig, Professorin für theoretische Informatik an der TU Kaiserslautern, zum Thema *Wie die Ethik in den Rechner kommt* vor. Sie ist Mitglied in verschiedenen Politikberatungsgremien, unter anderem im Koordinationsgremium des Netzwerks Verbraucherforschung und in der Enquete-Kommission Künstliche Intelligenz des Bundestages. Ihre Arbeit wurde mehrfach ausgezeichnet, unter anderem mit dem Communicator-Preis der DFG und des Stifterverbandes 2019. In ihrem Vortrag, der auch von ihrem neuen Buch *Ein Algorithmus hat kein Taktgefühl: Wo künstliche Intelligenz sich irrt, warum uns das betrifft und was wir dagegen tun können* inspiriert war, erklärte sie, an welchen Stellen beim Einsatz von Künstlicher Intelligenz die Antwort der Künstlichen Intelligenz durch menschliche Entscheidungen beeinflusst wird, und ermunterte die Anwesenden, sich bei solchen Entscheidungen einzumischen und sicherzustellen, dass diese ethisch getroffen werden.

Beide Veranstaltungen waren rege besucht und lösten bei den Anwesenden eine rege Anteilnahme an den an den Vortrag anschließenden Gesprächen aus.

In order to open its doors a bit further for the general public and the local region, Schloss Dagstuhl, together with the town of Wadern, initiated a new series of events: the *Dagstuhler Gespräche* (“Dagstuhl conversations”). The interested public will be introduced to a broad spectrum of topics from computer science, as well as to practical applications of those topics in everyday life or commercial processes. The talks are also meant to encourage the dialogue between decision makers and framers in industry and politics on the one hand and the interested public on the other hand.

The talk on May 17, 2019 was given by the former President of the Gesellschaft für Informatik e.V. (German Informatics Society) Prof. Dr. Hannes Federrath. Under the title “Was weiß das Internet über mich?” (What does the Internet know about me?), he explained what traces one leaves when on the Internet. He discussed how one can protect oneself from profiling and what data protection rights one has on the Internet.

One more event in this series took place: On November 24, Prof. Dr. Katharina Zweig, Professor for theoretical computer science at the TU Kaiserslautern, talked about the topic *Wie die Ethik in den Rechner kommt* (How ethics gets into the computer). She is a member of various political advisory bodies, inter alia the coordinating board of the network for consumer science and in the Bundestag’s committee of inquiry on Artificial Intelligence. Her work has won several awards, for example the communicator-award of the DFG (German Research Foundation) and the Stifterverband (Donors’ association for the promotion of humanities and sciences in Germany). In her talk, which was also inspired by her new book *Ein Algorithmus hat kein Taktgefühl: Wo künstliche Intelligenz sich irrt, warum uns das betrifft und was wir dagegen tun können* (An algorithm has no tact: Where artificial intelligence errs, why that matters to us, and what we can do about it), she explained where in the use of artificial intelligence its answers are influenced by human decisions. She encouraged the audience to take a hand in such decisions and make sure they are made ethically.

Both talks were well attended and the discussions afterwards were lively. The Dagstuhler Gespräche will certainly see a continuation in the next year.

[70] https://www.dagstuhl.de/19503

8 Einrichtungen
Facilities

Das Zentrum verfügt über drei Standorte; der Hauptstandort ist Schloss Dagstuhl in Wadern. Die Geschäftsstelle mit Sachbearbeitungsteam und wissenschaftlichen Mitarbeitern, die für die Dagstuhl-Seminare und Perspektiven-Workshops verantwortlich sind, befinden sich auf dem Campus der Universität des Saarlandes in Saarbrücken, während der Bibliographiedienst durch wissenschaftliche Mitarbeiter in Räumlichkeiten der Universität Trier betreut wird. Der Dagstuhl-Verlagsdienst befindet sich in Saarbrücken und Wadern.

The institution operates from three sites: the main site is Schloss Dagstuhl in Wadern. The administrative office and the scientific staff operating the Dagstuhl Seminars and Perspectives Workshops are located on the campus of Saarland University in Saarbrücken, while the scientific staff operating the Bibliographic Services are located in offices on the campus of the University of Trier. Dagstuhl Publishing is located in Saarbrücken and Wadern.

Hauptstandort in Wadern 8.1 *Main Site in Wadern*

Der Hauptstandort in Wadern umfasst das historische Schloss (gebaut um 1760) mit einem Anbau aus den 1970ern, einem 1993 fertiggestellten Erweiterungsbau, in dem sich Forschungsbibliothek, Hörsäle, Gästezimmer, Büros und Infrastruktur befinden, und ein 2012 fertiggestelltes Gästehaus mit Gästezimmern, einem Konferenzraum und Räumlichkeiten der Gebäudeverwaltung. Alle Einrichtungen in Wadern sind ganzjährig in Betrieb, abgesehen von je zwei Wochen im Sommer und Winter, die für größere Instandhaltungsarbeiten genutzt werden.

Die Kapazitäten von Dienstleistungen und Räumlichkeiten zur Veranstaltung von Seminaren sind genau aufeinander abgestimmt: Das Zentrum hat 71 Gästezimmer, davon sind 18 Doppelzimmer, sodass insgesamt 89 Teilnehmer übernachten können. Bei Normalbetrieb finden parallel zwei Seminare mit jeweils 30 und 45 Teilnehmern statt, wobei jedem Seminar ein Hörsaal für 35 bzw. 60 Personen zur Verfügung steht. Obwohl so eine Gesamtsumme von 75 Teilnehmern entsteht, ist es nur selten notwendig, Seminargäste in Doppelzimmern oder einem nahegelegenen Hotel unterzubringen. Die Obergrenze von 71 Zimmern wird regelmäßig erreicht, weshalb es wohl kaum Möglichkeiten gibt, die Nutzung unserer Einrichtungen weiter auszubauen.

The 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.

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.

Alle Hörsäle und andere Tagungsräume sind mit Tafeln und/oder Whiteboards augestattet.

Daneben gibt es über das ganze Zentrum verteilt weitere Räume, in denen Gäste sich in entspannter Atmosphäre treffen und diskutieren können. Insbesondere am Abend

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 is available for those who wish to go through their presentations without physical access to a computer.

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.

All lecture halls and meeting rooms are equipped with blackboards and/or whiteboards.

The center also offers a variety of other spaces where guests can sit and work together in a relaxed atmosphere. Particularly in the evening, guests gravitate towards the

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 Mitarbeitern der Küche frisch zubereitet. Der Schwerpunkt liegt dabei auf leichtem Essen während des Tages, um unsere Gäste nicht zu ermüden, und auf warmen Gerichten am Abend. Dies steht ein wenig im Widerspruch zur deutschen Tradition, kommt aber der Mehrheit der internationalen Gäste des Zentrums durchaus entgegen.

Sowohl die Zutaten als auch die Gerichte wechseln saisonal. An warmen Sommerabenden wird auf Anfrage auf der Terrasse vor dem Speisesaal gegrillt, unter anderem saarländische Schwenker, eine lokale Variante des Grillsteaks, die unter dauerndem Schwenken des Grillrostes zubereitet wird. In den kalten Monaten steht einmal wöchentlich ein schmackhafter Eintopf auf dem Speiseplan. Über das Jahr hinweg wird eine ausgewogene Mischung an regionalen und internationalen Spezialitäten aus neuen sowie bewährten und beliebten Rezepten angeboten. Im Allgemeinen sind die angebotenen Gerichte im Sommer etwas leichter und im Winter ein wenig schwerer. Die Küche arbeitet nach dem HACCP-Konzept (Hazard Analysis and Critical Points Concept) und hält sich an die Kennzeichnungspflicht von Allergenen, zu der alle lebensmittelverarbeitenden Betriebe verpflichtet sind. Des Weiteren achten wir auf deklarationsfreie Zusatz- und Konservierungsstoffe.

Alle Gäste, die aus medizinischen oder ethischen Gründen Einschränkungen bei der Speiseauswahl haben, können sich vor dem Seminar bei Schloss Dagstuhl melden. Unsere Küchenmitarbeiter erarbeiten gerne individuelle Lösungen für jeden Gast, soweit es irgend möglich ist. Gäste, die koscheres Essen benötigen, haben die Möglichkeit, mitgebrachte abgepackte Speisen selbst zu erhitzen.

Um unseren Gästen trotz eines begrenzten Budgets eine ausgewogene Qualität anbieten zu können, bietet unsere Küche ein Frühstücksbüffet, dienstags bis donnerstags abhängig von den personellen Kapazitäten ein Mittagsbuffet sowie ein Menü am Abend an. Montags und freitags wird aus logistischen Gründen auch am Mittag ein 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.

wine cellar and upstairs café, two of the coziest places in the house and great places for continuing a productive discussion in a comfortable atmosphere.

Dagstuhl's Kitchen

The dining experience at Dagstuhl is an important part of the center's scientific program. Seating arrangements are mixed deliberately in order to break up cliques and encourage guests to talk to as many different people as possible during the course of their stay. Large tables in the dining hall promote collaborative interaction during meals.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy, and tasty meals. Everything is freshly prepared each day by the kitchen's staff. The focus is on lighter fare during the day in order to aid scientists' concentration, and on a warm meal in the evening, breaking with the German tradition of a cold evening meal while matching the internationality of the center's guests.

Both ingredients and dishes vary with the seasons. On warm summer evenings, guests are invited on demand to partake of grilled *Schwenker* (the local variant of barbecued steak) on the outdoor patio adjacent to the dining hall. During the cold winter months, warm soups appear on the menu weekly. In general, the kitchen tries to keep meals lighter in the summertime and heavier in the winter, offering a blend of regional and international dishes year-round that include some new recipes and many tried-and-true Dagstuhl favorites. The kitchen works in accordance with the HACCP Concept (Hazard Analysis and Critical Points Concept) and adheres to the mandatory labeling of allergens, which is required of all food processing establishments. Food additives and conservatives for which labeling is non-mandatory are also carefully monitored.

All guests with special dietary requirements due to ethical or health reasons can announce their needs previous to the events. Our kitchen staff will then work out individual solutions if at all possible. Guests who need kosher meals can heat up ready-to-eat meals for themselves.

To accomplish all of this within a reasonable budget, the center offers a buffet-style breakfast and a set evening meal served by the kitchen's friendly and dedicated staff. From Tuesday to Thursday the kitchen offers a buffet-style lunch depending on the staff capacities. Due to logistical reasons, a set meal is served at lunch on Mondays and Fridays. The large dining-hall, seating up to 80 persons, opens onto the castle garden and patio, and offers a relaxed, familiar atmosphere.

Small and late-morning breaks punctuate the daily routine. During the small coffee break in the morning, hot drinks are served outside the lecture halls. During the longer coffee break in the afternoon, hot drinks together with freshly baked cake are served in the dining 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 the kiosk in front of the café. Bread and cheese is served in the café and the wine cellar every night.

Kleine und große Pausen unterbrechen auf angenehme Weise die tägliche Routine und anstrengenden Diskussionen. In der kleinen Kaffeepause am Vormittag stehen vor den Vortragsräumen heiße Getränke auf einem Kaffeewagen bereit. In der großen Kaffeepause am Nachmittag wird den Gästen im Speiseraum neben heißen Getränken auch frisch gebackener Kuchen angeboten. Darüber hinaus gibt es im Gästehaus, der „alten" Cafeteria und dem Weinkeller jeweils einen Kaffeevollautomaten zur Zubereitung von Kaffee, Kakao und Tee. Im Kiosk vor der Cafeteria können Gäste Snacks erwerben. Abends gibt es in der Cafeteria und im sogenannten Weinkeller einen Gruß aus der Küche, bestehend aus Brot und einer Käseauswahl.

Kinderbetreuung

Schloss Dagstuhl bietet Teilnehmern, die mit Kindern anreisen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser Service kann gegen ein geringes Entgelt im Voraus gebucht werden. Alternativ ist es Eltern auch möglich, eine Begleitperson zur Betreuung des Kindes oder der Kinder mitzubringen. Schloss Dagstuhl kommt für die Unterkunft und Verpflegung der Kinder auf. Wenn statt Inanspruchnahme der Kinderbetreuung von Schloss Dagstuhl eine Betreuungsperson mitreist, hat diese ebenfalls freien Aufenthalt.

Dagstuhls Angebot der Kinderbetreuung für Eltern wird weiterhin gut genutzt. Im Jahr 2019 wurden 27 Kinder durch eine Tagesmutter und 23 weitere durch Verwandte bzw. durch die Eltern selbst betreut. Insgesamt beherbergte Schloss Dagstuhl 50 Kinder von Teilnehmern an 25 Veranstaltungen während 20 Wochen.

Freizeit und Ambiente

Die Freizeitanlagen auf Schloss Dagstuhl wurden so gestaltet, dass sie auf unterschiedliche Art und Weise sowohl tagsüber als auch abends die Kommunikation zwischen den Seminarteilnehmern fördern. Die Mischung aus Arbeit und Freizeit in entspannter, familiärer Atmosphäre ist ein wichtiger Bestandteil des Dagstuhl-Konzepts. Gäste leben und arbeiten zusammen in einem Komplex aus drei Gebäuden, im Zentrum das historische Schloss, wo sie rund um die Uhr freien Zugang zu den zahlreichen Freizeiträumen und -anlagen haben. Musikalische Gäste können ihre Fertigkeiten im barocken Musiksaal zu Gehör bringen, wo ein Flügel und diverse andere Instrumente wie z. B. zwei Konzertgitarren zur Verfügung stehen. Unser Zentrum verfügt außerdem über eine Sauna, einen Billardtisch, Tischfußball, Mountainbikes, eine Dartscheibe, einen Freizeitraum mit Fitnessgeräten und Tischtennis sowie einen Außenbereich mit Volleyballnetz.

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 2019, Dagstuhl hosted 50 children, 27 of whom were cared for by a nanny on site, 23 by relatives or their parents. Participants of 25 events in 20 weeks were thus able to attend although they were traveling with their children.

Leisure Facilities

Leisure facilities at Schloss Dagstuhl are designed to encourage and support communication among seminar participants in different settings throughout the day and evening. This work/life continuum within a relaxed, informal setting is an important part of the Dagstuhl concept. Guests live and work together in a complex of three buildings, the historical manor house ("Schloss") in the middle, and enjoy full access to the center's many unique rooms and facilities around the clock. Musically talented guests are welcome to exercise their skills in the baroque music room on the upper floor of the historical main building, which features a grand piano and various other instruments, e.g., two concert guitars. Schloss Dagstuhl also has a full sauna, a pool table, table football facilities, mountain bikes, a dartboard, and a recreation room with gym equipment and table tennis as well as outdoor sports grounds featuring a volleyball net.

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 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 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 Mitarbeiter haben ihren Standort an der Universität Trier. Die Ende 2010 zunächst auf Basis zweier Projekte gestartete Zusammenarbeit zwischen Schloss Dagstuhl und der Universität Trier wurde im November 2018 in eine offizielle und permanente Außenstelle von Schloss Dagstuhl auf dem Campus der Universität Trier überführt. Dabei profitiert das dblp-Team von der engen Zusammenarbeit mit der Abteilung Informatikwissenschaften und als externer Partner im Digital Research and Bibliographic Meta Data Lab des Center for Informatics Research and Technology (CIRT).

8.3 *Dagstuhl at University of Trier*

The scientific and editorial staff working on the *dblp computer science bibliography* is located at the Dagstuhl offices at the University of Trier. Initially based on a project-based cooperation between Schloss Dagstuhl and the University of Trier which was first established in 2010, in November 2018, an official and permanent Schloss Dagstuhl branch office has been established on the campus of the University of Trier. In Trier, the dblp team benefits from the close cooperation with the University's department of computer sciences, and as an external partner in the Center for Informatics Research and Technology (CIRT) lab for Digital Research and Bibliographic Meta Data.

9 Zentrale Dienste
Central Services

Schloss Dagstuhl verfügt über zwei zentrale Dienste: die IT-Abteilung und eine Forschungsbibliothek. Beide Einrichtungen befinden sich am Hauptstandort in Wadern.

Schloss Dagstuhl has two central services: the IT service and a research library, which are both located at the main site in Wadern.

Bibliothek 9.1 *Research Library*

Zur wissenschaftlichen Literatur- und Informationsversorgung der Seminarteilnehmer unterhält Schloss Dagstuhl eine hervorragende Forschungsbibliothek für Informatik.

Die Bibliothek ist für Wissenschaftler vor Ort rund um die Uhr und für externe Wissenschaftler nach Absprache zugänglich. Zur digitalen Informationsinfrastruktur gehören ein Online-Bibliothekskatalog, ein modernes Discovery-System zur Artikelrecherche sowie zahlreiche Angebote für den Online-Zugriff auf wissenschaftliche Publikationen.

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. Zur Optimierung der Autorenidentifikation werden die ORCIDs der Personennamen im Bibliothekskatalog erfasst.

Außerdem wird der Name eines jeden Seminarteilnehmers in der Online-Teilnehmerliste mit seinen oder ihren in der dblp-Literaturdatenbank erfassten Veröffentlichungen verlinkt. Diese Maßnahmen ermöglichen den Seminarteilnehmern einfachen und schnellen Zugriff auf seminarrelevante Literatur.

Die Bibliothek verfügt über einen umfangreichen Buchbestand, der Zugriff auf aktuelle Forschungspublikationen wie Konferenzbände und wissenschaftliche Zeitschriften erfolgt ausschließlich digital.

- Der Buchbestand orientiert sich am wissenschaftlichen Seminarprogramm. 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 35 000 Informatikbücher. Die Metadaten werden standardisiert erfasst und mit Hyperlinks angereichert, die durch persistente Adressierung (DOIs) verlässlich verlinkt sind.
- 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 teilweise 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.
- Wissenschaftliche Fachzeitschriften sind eine wesentliche Voraussetzung für exzellente Forschung. Häufig werden in Zeitschriften erweiterte Fassungen von

Schloss Dagstuhl maintains an excellent research library for computer science to provide seminar participants with scientific literature and information.

The library is accessible to on-site researchers around the clock and to external researchers by appointment. The digital information infrastructure includes an online library catalog, a modern discovery system for article research as well as numerous options for online access to scientific publications.

For each seminar, an individual book exhibition is compiled, consisting of books written or edited by seminar participants. The authors who are present at the seminar are asked to sign their own books. In order to optimize the author identification, the ORCIDs of the authors' names are recorded in the library catalog.

In addition, the name of each seminar participant will be linked in the online list of participants with their publications recorded in the dblp literature database. These measures provide seminar participants with easy and quick access to the literature relevant to the seminar.

The library maintains an extensive collection of books. Access to current research publications such as conference proceedings and scientific journals is exclusively digital.

- The book collection is oriented towards the scientific seminar program. New acquisitions focus on books which are related to Dagstuhl Seminars and Perspectives Workshops or which were written by seminar organizers or participants. In addition, the library receives numerous books as donations from publishers and authors. Currently, the library has about 35,000 books on computer science. The metadata are recorded in a standardized way and enriched with hyperlinks, which are reliably linked by permanent addressing (DOIs).
- Contributions in conference proceedings represent the most important part of the literature in computer science. The library has subscribed to the complete ACM and IEEE proceedings electronically. Earlier volumes are also partly available in printed form. The Springer-Nature publishing group donates all volumes of the series Lecture Notes in Computer Science (LNCS) to the library both in print and in electronic form. The library thus has print copies of all published volumes from volume 1 onwards.
- Scientific journals are essential for excellent research. Journals often publish extended versions of results that were previously published in conference proceedings. The library provides access to several thousand digital scientific journals. Most of them are included in journal packages licensed in cooperation with nationwide consortia, such as DFG-funded national and alliance

Ergebnissen veröffentlicht, die zuvor in Konferenzbänden publiziert wurden. Die Bibliothek bietet Zugriff auf mehrere Tausend digitale Fachzeitschriften. Die meisten sind in Zeitschriftenpaketen enthalten, die in Kooperation mit deutschlandweiten Konsortien lizenziert sind, beispielsweise DFG-geförderte National- und Allianzlizenzen sowie von der Leibniz-Gemeinschaft geförderte Konsortiallizenzen.

- Die Bibliothek ermöglicht den benutzerfreundlichen Online-Zugriff auf über 7 000 deutschlandweite und internationale Zeitungen und Magazine aus über 120 Ländern.

licenses as well as consortium licenses funded by the Leibniz Association.

- The library enables user-friendly online access to over 7,000 Germany-wide and international newspapers and magazines from over 120 countries.

Zusammenarbeit

Schloss Dagstuhls Forschungsbibliothek ist mit zahlreichen überregionalen Bibliotheksdatenbanken vernetzt. Der komplette Zeitschriftenbestand ist in der Zeitschriftendatenbank (ZDB) nachgewiesen. Zusätzlich ist der Bestand an elektronischen Zeitschriften in der kooperativen bundesweiten Elektronischen Zeitschriftenbibliothek (EZB) erfasst. Darüber hinaus wird der aktuelle Buchbestand im K10plus, der gemeinsamen Katalogisierungsdatenbank von GBV und SWB mit über 180 Millionen Nachweisen, nachgewiesen.

Diese Datenbanken bilden die Grundlage für den deutschlandweiten und internationalen Leihverkehr der Bibliotheken. Somit steht der Zeitschriftenbestand auch standortübergreifend und überregional für Fernleihzwecke zur Verfügung.

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.

Die Bibliothek war in den letzten Jahren Mitglied bei LITexpress, der virtuellen Bibliothek für Rheinland-Pfalz, das Saarland und die deutschsprachige Gemeinschaft in Belgien, ein Medienverleihservice für die Einwohner dieser Regionen. LITexpress wurde jedoch zum 31.12.2018 eingestellt.

Schloss Dagstuhls Fachbibliothek ist institutionelles Mitglied des Deutschen Bibliotheksverbandes (DBV). Die Bibliothekarin Frau Meyer ist persönliches Mitglied im Berufsverband Information Bibliothek e.V. (BIB).

Collaboration

Schloss Dagstuhl's research library is connected to numerous national library databases. The complete journal inventory is recorded in the Zeitschriftendatenbank (ZDB). In addition, the inventory of electronic journals is recorded in the cooperative nationwide Electronic Journals Library (Elektronische Zeitschriftenbibliothek, EZB). Furthermore, the current book stock is recorded in K10plus, the joint cataloging database of GBV and SWB with over 180 million records.

These databases form the foundation for the libraries' nationwide and international lending system. Thus the journal collections are also available for inter-library loan purposes across locations and regions.

There is also a close cooperation between Schloss Dagstuhl and the Saarland University and State Library (SULB), the Campus Library for Computer Science and Applied Mathematics at Saarland University and the library of the Leibniz Institute for New Materials (INM), all of which are located in Saarbrücken.

In recent years, the library has been a member of LITexpress, the virtual library for Rhineland-Palatinate, Saarland, and the German-speaking community in Belgium, a media lending service for the inhabitants of these regions. However, LITexpress was discontinued on 31.12.2018.

Schloss Dagstuhl's specialized library is an institutional member of the German Library Association (Deutscher Bibliotheksverband, DBV). The librarian Ms. Meyer is a personal member of the Professional Association Information and Libraries (Berufsverband Information Bibliothek e.V., BIB).

Spenden an die Bibliothek

Die Bibliothek von Schloss Dagstuhl profitiert von zahlreiche Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2019 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 672 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. In 2019, 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 672 volumes during the year 2019 as donations from publishing houses and seminar participants.

IT-Service 9.2 IT Service

Die IT-Abteilung bietet umfassenden Support für alle internen Vorgänge an den 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 375 Mbit/s-Leitungen sichergestellt, betrieben durch den DFN e.V. (Deutsches Forschungsnetz).
- Fahrbare ebenso wie fest montierte Präsentationsmöglichkeiten in den Tagungsräumen. In den größeren Tagungsräumen können Vortragende den vorhandenen oder den eigenen Laptop verwenden.
- Zugang zu Netzwerkfarbdruckern, einem Scanner und einem Kopierer.
- Zugang zu gemeinschaftlich genutzten Computern mit den Betriebssystemen Microsoft Windows, Apple Mac OS X und Linux.
- Technischen Support für Seminarteilnehmer und Mitarbeiter von Schloss Dagstuhl.

Der IT-Service verwaltet (virtuelle) Server für alle Abteilungen, z.B.

- einen Webserver, auf dem sich Schloss Dagstuhls Internetpräsenz befindet (https://www.dagstuhl.de), die Informationen für Teilnehmer, zum Seminarprogramm usw. enthält,
- einen Server, auf dem sich DROPS befindet, Schloss Dagstuhls Publikationsplattform (http://drops.dagstuhl.de),
- den dblp-Server (https://dblp.dagstuhl.de und https://dblp.org).

Darüber hinaus stellt der IT-Service Tools für das gemeinschaftliche Arbeitsumfelds zur Verfügung und hält sie in Stand, z.B. *Sihot* (eine Software zur Organisation von Gastdaten), MySQL-Datenbanken, ownCloud (ein Cloud-basiertes Speichersystem) und weitere.

The IT service provides comprehensive support for all internal operations at all three sites. Moreover, it provides IT infrastructure, services, and support for all guests of Dagstuhl events.

This service includes – among others – the following:

- Internet access via Ethernet and Wi-Fi throughout all rooms. For Wi-Fi access Schloss Dagstuhl offers personal accounts and also takes part in the *eduroam* service[71] (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 375 Mbit/s 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, ownCloud (a cloud-based storage system), and several others.

[71] *eduroam* (education roaming) is a world-wide roaming access service developed for the international research and education community, see https://www.eduroam.org.

SIAM – Society for Industrial and Applied Mathematics
http://www.siam.org

Springer-Verlag GmbH | Springer Science+Business Media
http://www.springer.com

Fig. 9.1
Donations from publishers to the Dagstuhl library.

10 Kunst
Art

Dagstuhl als Galerie 10.1 *Dagstuhl as Art Gallery*

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des 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 vier Ausstellungen (siehe Fig. 10.1), die im Jahr 2019 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 has continued to supervise the Schloss Dagstuhl art exhibitions following his retirement as the center's Scientific Director in April 2014. The center holds approximately three to four art exhibits per year, with each exhibit generally running for two to three months.

Until now, the exhibitions were organized by artists and individual collectors. The year 2016, however, saw the establishment of a cooperation between Saarland-Sporttoto GmbH (Saartoto for short), Hochschule für Bildende Künste Saar (university of art and design; HBKsaar for short), and Schloss Dagstuhl, which makes 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 the 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 four exhibitions (cf. Fig. 10.1) hosted by Schloss Dagstuhl in 2019 are described below. Current exhibitions are open to the interested public upon request.

»ELLIPSE – 1999-2019« Works from Sabrina Sperl | January 14 to April 5, 2019
»Malerei auf Leinwand und Papier 2017–2019« Works from Werner Constroffer | April 29 to July 19, 2019
»Solutions« Works from Gisela Zimmermann | August 19 to October 11, 2019
»Lost Places« Works from Winfried Groke | October 21 to December 20, 2019

Fig. 10.1
Art exhibitions in 2019.

»ELLIPSE – 1999–2019«

Am Anfang war der Bottich aus dem Baumarkt. Schwarz. Aus Plastik. Ausgehend von diesem Gegenstand läuft der Prozess seiner künstlerischen Transformation seit 20 Jahren. Geblieben ist die Form der Ellipse, die zugleich namengebend für die Werkgruppe steht. Sie erscheint in den Arbeiten als Leerstelle oder beschreibt fragmentarische räumliche Zusammenhänge. Die gezeigten Werke stehen in ihrer formalen Verschiedenheit exemplarisch für die Arbeitsweise von Sabrina Sperl, geboren im 1969 in Wuppertal. Rückbezüge suchend und einem roten Faden folgend, setzt sie die Ellipse seit Beginn ihres Studiums 1998 als universelle Form in ihrer malerischen Position ein.

»Malerei auf Leinwand und Papier 2017–2019«

Werner Constroffer, 1949 in Saarlouis geboren, hat von 1970 bis 1974 an der Werkkunstschule Saarbrücken bei Oskar Holweck und Robert Sessler studiert. Er arbeitete als Grafikdesigner in verschiedenen Werbeagenturen und ist seit 1983 mit seinen Arbeiten in Kunstausstellungen vertreten. Er ist Mitglied des Saarländischen Künstlerbund.

Durch die Fotografie, das Aufnehmen und Bewahren des Augenblicks, schafft er einen Fundus von real Gelebtem. Die fotografischen Lebensbilder dienen ihm als Vorlage, die er zeichnerisch bearbeite. Erzählerisch beschreibt die Linienkontur Gelebtes. Sie belebt das vordergründig Reale und schafft Spielraum für den Künstler und Betrachter.

»Solutions «

Die Malerei von Gisela Zimmermann, geboren im 1965 in Merzig an der Saar, ist geprägt von dynamischen Prozessen im Akt der Entstehung, die der abschließenden Bildgestalt jene kraftvoll-energetische Aufladung verleihen, die den Betrachter in die expressiv gestalteten Farbräume hineinzieht.

In den neueren Arbeiten finden sich eigentümlich-geometrisierende Strukturen als Farbakzente in magischen Dunkelräumen wie auch gestisch-eruptive Entladungen einer koloristischen Spannung als intuitive Improvisation der Bildgestaltung. Die hohe Suggestivkraft der Bilder zieht Betrachter emotionalisierend in den Bann und fesselt mit einer kaum vergleichbaren Wirkmächtigkeit und atmosphärischen Intensität.

»Lost Places«

Wo sonst kann man eine Handlung so dicht und auf das Wesentliche reduziert erleben als im Theater? Winfried Groke, geboren im 1952 in Menden Sauerland, hat sich manche Tugend des Theaters zu eigen gemacht. Anders als auf den Brettern, die die Welt bedeuten, besetzt der

»ELLIPSE – 1999–2019«

In the beginning, there was the tub from the hardware store. Black. Plastic. On the basis of this object, a process of artistic transformation has been going on for 20 years. What has remained is the form of the ellipse, which also gives this collection of works its name. It appears in the works as a blank space or describes fragmentary spatial contexts. In their formal diversity, the works on display are exemplary for the working method of Sabrina Sperl, born in 1969 in Wuppertal. Searching for references and following a red thread, she has used the ellipse as a universal form in her artistic position since the beginning of her studies in 1998.

»Malerei auf Leinwand und Papier 2017–2019«[72]

Werner Constroffer, born in 1949 in Saarlouis, has studied at the Werkkunstschule Saarbrücken with Oskar Holweck and Robert Sessler between 1970 and 1974. He has worked as a graphics designer in several advertising agencies and, since 1983, his work is featured in art exhibitions. He is a member of the Saarland Artists Association.

Through photography, the capturing and preserving of the moment, he creates a wealth of real life. The photographic images of life serve him as a template, which he elaborates in drawings. In a narrative way, the line contour describes the reality that has been lived. It animates the ostensibly real and creates space for the artist and the viewer.

»Solutions«

The painting of Gisela Zimmermann, born in 1965 in Merzig at the Saar, is characterized by dynamic processes in the act of creation, which give the final picture-Gestalt the powerful energetic charge which draws the viewer into the expressively designed color spaces.

In the more recent works, one finds peculiar geometricizing structures as color accents in magical dark spaces as well as gestural-eruptive discharges of a coloristic tension as an intuitive improvisation of the picture design. The high suggestive power of the pictures draws the viewer under an emotional spell and captivates them with a unique power and atmospheric intensity.

»Lost Places«

Where else can one experience a narrative so dense and reduced to the essential as in the theatre? Winfried Groke, born in 1952 in Menden in Sauerland, has made many virtues of theatre his own. Unlike on the stage, the photographer does not cast actors in the roles, but objects.

[72] engl. Paintings on Canvas and Paper 2017-2019

Fotograf die Rollen nicht mit Schauspielern, sondern mit Objekten. In seinem „Theater der Gegenstände“ erhalten beispielsweise Messer und Gabel anthropomorphe Züge, wie umgekehrt das menschliche Antlitz verdinglicht und häufig in Form von Büsten erscheint. Grokes Fotografien scheinen Momente einer zwischen den Gegenständen sich abspielenden Handlung festzuhalten.

In his "Theater of Objects", for example, knives and forks take on anthropomorphic features, just as the human face is conversely objectified and often appears in the form of busts. Groke's photographs seem to capture moments of a story taking place between the objects.

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 2019 erhielt Schloss Dagstuhl insgesamt 5 083 € von verschiedenen Spendern. Wir möchten diese Stelle nutzen, allen Spendern, die 2019 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

Nähere Informationen und aktuelle Neuigkeiten finden sich auf der Kunst-Webseite[73] 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 2019, Schloss Dagstuhl received a total of 5,083 € from various donors. We would like to thank all donors who contributed to Dagstuhl's art collection in 2019.

For further information and current news about Dagstuhl's art program, please visit Dagstuhls's art webpage[73].

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 25 Jahren kamen so ungefähr 180 Kunstwerke zusammen. Auch durch diese Initiative angeregt und verstärkt erhielt Dagstuhl in den vergangenen Jahren weitere Spenden von Künstlern und Mäzenen. Die Arbeiten kommen in den Räumen des Zentrums in Wadern sowie in der Geschäftsstelle in Saarbrücken sehr gut zur Geltung.

The art collection, continually praised by guests, was initiated by Founding Director Professor Wilhelm. It was his idea to use works of art in order to enliven the New Building as well as the dining room opened in 1994 and 1995, respectively. To this end, Professor Wilhelm launched the exhibitions described above. Assisted by the artists, one picture from each exhibition was chosen and donors were drummed up. Thus, approximately 180 works of art could be acquired over the last 25 years. Additionally, this initiative has increasingly encouraged artists and patrons to make donations. All of the pictures adorn the rooms of Schloss Dagstuhl in Wadern as well as the Dagstuhl Office in Saarbrücken.

[73] https://www.dagstuhl.de/art/

11 Struktur der Gesellschaft
Structure of the Company

Gründung und Gesellschafter 11.1 *Formation and Shareholders*

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit elf Gesellschaftern (siehe Fig. 11.1) organisiert. Dies sind einerseits die vier Gesellschafter, die Schloss Dagstuhl gegründet haben, nämlich die Gesellschaft für Informatik e. V. (GI), die Universität des Saarlandes, die Technische Universität Kaiserslautern und das Karlsruher Institut für Technologie (KIT). Als vier weitere Gesellschafter wurden 1994 die Technische Universität Darmstadt, die Johann Wolfgang Goethe-Universität Frankfurt am Main, die Universität Stuttgart und die Universität Trier aufgenommen. Drei international renommierte Forschungsinstitute, das Institut National de Recherche en Informatique et en Automatique (INRIA, Frankreich), das Centrum Wiskunde & Informatica (CWI, Niederlande) und die Max-Planck-Gesellschaft (MPG, Deutschland) wurden 2005/2006 als weitere Gesellschafter aufgenommen.

Aufgrund eines Beschlusses der Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung (heute Gemeinsame Wissenschaftskonferenz) wurde das Zentrum mit Wirkung zum 1. Januar 2006 als Serviceeinrichtung für die Forschung in die gemeinsame Forschungsförderung von Bund und Ländern aufgenommen. Es ist seit 2005 Mitglied der Leibniz-Gemeinschaft. Entsprechend wurde 2008 der Name des Zentrums von vormals „Internationales Begegnungs- und Forschungszentrum für Informatik“ in „Schloss Dagstuhl – Leibniz-Zentrum für Informatik“ geändert.

Schloss Dagstuhl wurde im Juli 2009 erstmals durch die Leibniz-Gemeinschaft evaluiert. Die Stellungnahme der Evaluierungs-Kommission vom März 2010 war sehr positiv: Schloss Dagstuhl widme sich mit herausragendem Erfolg seiner Aufgabe, die internationale Informatikforschung mit einem Seminarzentrum für wissenschaftliche Veranstaltungen zu unterstützen. Schloss Dagstuhl wurde 2016 erneut mit hervorragendem Ergebnis evaluiert. In der Stellungnahme des Senats der Leibniz-Gemeinschaft wurde das Veranstaltungsprogramm und die Beteiligung an der Literaturdatenbank dblp als „exzellent“ bewertet, während der Bereich Open Access (Publishing) als „sehr gut“ bewertet wurde.

Schloss Dagstuhl is operated as a non-profit organization by eleven associates (cf. Fig. 11.1), including its four founding associates: the Gesellschaft für Informatik e. V.[74] (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[75] (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”[76] to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”[77] 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.

[74] engl.: German Informatics Society
[75] engl.: Federal/State Government Commission for Educational Planning and Research Promotion
[76] engl.: International Conference and Research Center for Computer Science
[77] 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.

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee consisting of at least the chairperson of the Supervisory Board and the chairperson of the Scientific Advisory Board. The term of office of the Scientific Director is five years.

The Technical Administrative Director is responsible for technical and administrative tasks. Since July 2014, Ms Heike Meißner holds this position.

Gremien der Gesellschaft 11.3 *Dagstuhl Bodies*

Die Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH werden durch drei Gremien unterstützt. Es sind die folgenden:

- Wissenschaftliches Direktorium
- Wissenschaftlicher Beirat
- Kuratorium

Details zu den Gremien werden in den folgenden Abschnitten ausgeführt.

The organs of Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH are supported by the following bodies:

- Scientific Directorate
- Scientific Advisory Board
- Industrial Curatory Board

Detailed information about these boards can be found in the sections below.

Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 11.4) ist für die Realisierung des Gesellschaftszwecks in fachlich-wissenschaftlicher Hinsicht verantwortlich. Es hat das Forschungs- und Veranstaltungsprogramm der Gesellschaft festzulegen, seine fachlich-wissenschaftliche Qualität zu sichern und seine Durchführung zu überwachen. Als wesentlicher Bestandteil dieser Aufgabe werden die Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von Mitgliedern des Wissenschaftlichen Direktoriums begutachtet. Auf den zweimal im Jahr stattfindenden Direktoriumssitzungen werden die Anträge diskutiert und es wird über ihre Annahme entschieden.

Der Wissenschaftliche Direktor gehört dem Wissenschaftlichen Direktorium an. Er empfiehlt dem Aufsichtsrat die Größe des Direktoriums. Neben den Gesellschaftern können das bestehende Wissenschaftliche Direktorium sowie der Beirat Kandidaten für das Wissenschaftliche Direktorium benennen. Die Auswahl der Kandidaten, die dem Aufsichtsrat zur Ernennung vorgeschlagen werden, obliegt dem Beirat zusammen mit dem Wissenschaftlichen Direktor.

Die Amtszeit der Mitglieder des Wissenschaftlichen Direktoriums – mit Ausnahme der des Wissenschaftlichen Direktors – beträgt drei Jahre. Sie beginnt am 1. November des Jahres ihrer Berufung und endet drei Jahre später am 31. Oktober. Wiederberufung ist möglich.

Scientific Directorate

The Scientific Directorate (see Fig. 11.4) is responsible for carrying out the company objectives from a technical and scientific point of view. It must determine the research and event program, ensure its technical and scientific quality, and monitor its execution. As a main task in support of this objective, members of the Scientific Directorate review proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In its 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.

Der Wissenschaftliche Beirat

Die Aufgaben des Wissenschaftlichen Beirats (siehe Fig. 11.5) werden nicht nur durch den Gesellschaftsvertrag festgelegt, sondern auch durch die Empfehlungen der Leibniz-Gemeinschaft. Im Sinne dieser wirkt der Wissen-

Scientific Advisory Board

The tasks of the Scientific Advisory Board (see Fig. 11.5) are not only defined by the Shareholders' Agreement, but also by the recommendations of the Leibniz Association. The latter stipulates two different ways in

schaftliche Beirat auf zwei Wegen bei der Qualitätssicherung mit. Zum einen berät er die Leitung in Fragen der Forschungs- und Entwicklungsplanung, nimmt Stellung zu den Programmbudgets und gibt Empfehlungen zum Ressourceneinsatz. Er unterstützt weiterhin den Aufsichtsrat bei wichtigen Entscheidungen zur Weiterentwicklung von Schloss Dagstuhl und bei der Gewinnung von Leitungspersonal. Zum anderen führt der Wissenschaftliche Beirat mindestens einmal zwischen je zwei Evaluierungen durch den Senatsausschuss Evaluierung (SAE) der Leibniz-Gemeinschaft ein Audit durch, bei dem die gesamte Einrichtung begutachtet wird. Ein Bericht über das Audit wird der Leitung, dem Aufsichtsrat und dem Senatsausschuss vorgelegt.

Der Wissenschaftliche Beirat sollte aus sechs bis zwölf international angesehenen, im Berufsleben stehenden Wissenschaftlern aus dem In- und Ausland bestehen. Die Amtszeit der Mitglieder beträgt vier Jahre, eine einmalige Wiederberufung ist möglich. Der Beirat wählt aus seiner Mitte einen Vorsitzenden. Der Wissenschaftliche Beirat tagt einmal im Jahr. Mitglieder des Beirats werden vom Aufsichtsrat auf Vorschlag des Beirats ernannt.

Das Kuratorium

Das Kuratorium (siehe Fig. 11.6) erfüllt eine Transmissionsfunktion zwischen Schloss Dagstuhl und den Forschungsabteilungen und Entwicklungslaboren der Industrie. Es hat die Aufgabe, die Akzeptanz des Zentrums in Verwaltung, Industrie und Wirtschaft abzusichern und als Förderungsorganisation die wirtschaftliche Basis des Zentrums zu verbreitern. Mitglieder des Kuratoriums werden vom Aufsichtsrat ernannt.

Nach seiner Geschäftsordnung hat das Kuratorium mindestens fünf Mitglieder, deren Amtszeit vier Jahre beträgt. Eine einmalige Wiederberufung ist möglich. Die Mitglieder des Kuratoriums unterstützen das Zentrum dabei, aktuelle Themen zu identifizieren und dazu geeignete zugkräftige Organisatoren aus der Industrie zu gewinnen. Sie werden ebenso gebeten, geeignete Personen aus der Industrie als Teilnehmer von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops zu benennen. Das industrielle Kuratorium tagt einmal im Jahr zusammen mit dem Wissenschaftlichen Beirat.

which the Scientific Advisory Board is involved in quality assurance. On the one hand, the board offers advice to the management with regard to research as well as development planning and issues comments on the program budget draft, making recommendations on the use of resources. It also assists the Supervisory Board in important decisions with regard to future development of the institute as well as the acquisition of management staff. On the other hand, it carries out an audit of the entire institute between two evaluations by the Senatsausschuss Evaluierung (SAE, Senate Committee Evaluation) of the Leibniz Association. A report on this audit is sent to the management, the Supervisory Board, and the SAE.

The Scientific Advisory Board should consist of six to twelve internationally reputable, well established scientists and academics from Germany and abroad. The term of office for members is four years and can be prolonged once. The Scientific Advisory Board members elect a chairperson from their midst. The board convenes once a year. Members are appointed by the Supervisory Board in accordance with the suggestions of the Scientific Advisory Board.

Industrial Curatory Board

The Industrial Curatory Board (see Fig. 11.6) performs a transmissional function between the center and the industrial R&D departments and laboratories. Its role is to secure acceptance of Schloss Dagstuhl within the business, industry and administrative communities, and as a promotional organization to broaden the economic basis of the center. Board members are appointed by the Supervisory Board.

According to its rules of procedure, the Industrial Curatory Board consists of at least five members whose term of office is four years. A one-off reappointment for a second term is possible. The board members help the center to identify current R&D topics for seminars and locate attractive organizers in industry. The Industrial Curatory Board is regularly called upon to propose suitable participants for Dagstuhl Seminars and Dagstuhl Perspectives Workshops known to it from its activities. It convenes once a year together with the Scientific Advisory Board.

Gesellschafter | Associates
Centrum Wiskunde & Informatica (CWI), The Netherlands
Gesellschaft für Informatik e. V., Germany
Institut National de Recherche en Informatique et en Automatique (INRIA), France
Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany
Karlsruher Institut für Technologie (KIT), Germany
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Berlin, Germany
Technische Universität Darmstadt, Germany
Technische Universität Kaiserslautern, Germany
Universität des Saarlandes, Germany
Universität Stuttgart, Germany
Universität Trier, Germany

Fig. 11.1
Associates.

Aufsichtsrat | Supervisory Board
Dr. Marc Brüser Ministerium für Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of Rhineland-Palatinate state
Prof. Dr.-Ing. Hannes Federrath Universität Hamburg, Germany | Representative of Gesellschaft für Informatik e. V.
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. Volker Lindenstruth Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany | Representative of Johann Wolfgang Goethe-Universität Frankfurt am Main
Dr. Svenja Marx Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government | *tenure started in October 2019*
Dr. Rainer Müssner Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government | *tenure ended in July 2019*
Prof. Dr. Arnd Poetzsch-Heffter Technische Universität Kaiserslautern, Germany | Representative of Technische Universität Kaiserslautern
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany | Representative of the Saarland
Prof. Dr. Ralph Schenkel Universität Trier | Representative of Universität Trier
Prof. Dr. Manfred J. Schmitt Universität des Saarlandes, Saarbrücken, Germany | Representative of Universität des Saarlandes
Prof. Dr. Peter H. Schmitt Karlsruher Institut für Technologie, Germany | Representative of Karlsruher Institut für Technologie
Prof. em. Dr.-Ing. Dr.-Ing. h. c. Roland Vollmar Karlsruher Institut für Technologie, Germany | Representative of Gesellschaft für Informatik e. V.
Cornelia Winter Gesellschaft für Informatik e. V., Bonn, Germany | Representative of Gesellschaft für Informatik e. V

Fig. 11.2
Supervisory Board members.

Geschäftsführung | Management
Heike Meißner (Technisch-administrative Geschäftsführerin | Technical Administrative Director) Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Wadern, Germany
Prof. Raimund Seidel, Ph. D. (Wissenschaftlicher Direktor | Scientific Director) Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Wadern and Universität des Saarlandes, Saarbrücken, Germany

Fig. 11.3
Management.

Wissenschaftliches Direktorium | Scientific Directorate
Prof. Dr. Elisabeth André Universität Augsburg, Germany | *tenure started in November 2019*
Prof. Dr.-Ing. Franz Baader TU Dresden, Germany | *tenure started in November 2019*
Prof. Gilles Barthe, Ph. D. IMDEA Software Institue, Madrid, Spain
Prof. Dr. Bernd Becker Albert-Ludwigs-Universität Freiburg, Germany | *tenure ended in October 2019*
Prof. Dr. Daniel Cremers Technische Universität München, Germany
Prof. Dr. Stefan Diehl Universität Trier, Germany | *tenure ended in October 2019*
Prof. Dr. Reiner Hähnle TU Darmstadt, Germany
Prof. Dr. Barbara Hammer Universität Bielefeld, Germany | *tenure started in November 2019*
Prof. Dr. Lynda Hardman Centrum Wiskunde & Informatica (CWI), Amsterdam and University of Utrecht, The Netherlands
Prof. Dr.-Ing. Oliver Kohlbacher Eberhard Karls Universität Tübingen, Germany
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany
Prof. Dr. Bernhard Nebel Albert-Ludwigs-Universität Freiburg, Germany | *tenure ended in October 2019*
Prof. Dr. Albrecht Schmidt Ludwig-Maximilians Universität München, Germany
Prof. Dr.-Ing. Wolfgang Schröder-Preikschat Friedrich-Alexander-Universität Erlangen-Nürnberg, 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
Prof. Dr. Heike Wehrheim Universität Paderborn, Germany
Prof. Dr. Verena Wolf Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Martina Zitterbart Karlsruher Institut für Technologie, Germany

Fig. 11.4
Scientific Directorate.

Wissenschaftlicher Beirat | Scientific Advisory Board
Prof. Dr. Christel Baier Technische Universität Dresden, Germany
Prof. Dr. Anja Feldmann Max-Planck-Institut für Informatik, Saarbrücken, Germany
Prof. Dr. Manuel V. Hermenegildo IMDEA Software Institute, Madrid and Technical University of Madrid, Spain
Prof. Dr. Ir. Joost-Pieter Katoen, PDEng RWTH Aachen, Germany | *tenure started in January 2019*
Prof. Dr. Claude Kirchner Institut National de Recherche en Informatique et en Automatique (INRIA), Villers-lès-Nancy, France | *tenure ended in December 2019*
Prof. Dr. Friedhelm Meyer auf der Heide Heinz Nixdorf Institute, Paderborn and Universität Paderborn, Germany | Chairman of the Scientific Advisory Board
Prof. Dr.-Ing. Dr. h. c. Andreas Reuter Heidelberg Laureate Forum Foundation, Heidelberg, Germany | *tenure ended in December 2019*
Prof. em. Dr. Dr. h. c. Otto Spaniol RWTH Aachen, Germany | *tenure ended in Decmeber 2019*

Fig. 11.5
Scientific Advisory Board.

Kuratorium | Industrial Curatory Board
Dr. Udo Bub Eötvös Loránd University, Budapest, Hungary | *tenure ended in December 2019*
Dr.-Ing. Uwe Franke Daimler AG, Böblingen, Germany | *tenure ended in December 2019*
Dr. Goetz Graefe Google, Madison, Wisconsin, United States
Dr. Tim Harris Amazon, Cambridge, United Kingdom
Dr. Michael May Siemens AG, München, Germany | *tenure ended in December 2019*
Dr.-Ing. Andreas Wierse SICOS BW GmbH, Stuttgart, Germany
Dr. Thomas Ziegert SAP SE, Darmstadt, 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 unter der strategischen Aufsicht eines Stiftungsrates (siehe Fig. 12.1). Der Verein wird von einem Vorstand (siehe Fig. 12.2 und Fig. 12.3) geleitet.

Das Stiftungsvermögen wird mit Hilfe einer professionellen und auf Stiftungen spezialisierten Vermögensverwaltungsgesellschaft sicher angelegt. Der Verein hat die ersten Jahre seiner Existenz dazu genutzt einen finanziellen Grundstock aufzubauen, und beginnt nun damit, die ersten Maßnahmen zur Förderung von Schloss Dagstuhl in die Wege zu leiten:

- Da in den vergangenen Jahrzehnten Schloss Dagstuhl auch durch die Verbindung von Kunst und Informatik bekannt geworden ist, und dadurch auch für die Forscher vor Ort stets besonders inspirierend ist, hat der Verein beschlossen, den Kunstbetrieb möglichst dauerhaft zu unterstützen und als ersten Schritt den Ankauf zweier Werke von Künstlern aus der Region finanziert. Darüber hinaus stehen Mittel zur fachlichen Sichtung und Archivierung des Kunstbestandes sowie gfls. die Ausarbeitung eines Vorschlages für die Hängung der Werke bereit.
- Eine weitere Maßnahme wurde durch die großzügige Zuwendung aus dem Nachlass eines verstorbenen Freundes von Schloss Dagstuhl ermöglicht: Da die Anbindung von Schloss Dagstuhl über den ÖPNV an den Bahnhof Türkismühle für unsere Tagungsteilnehmer oftmals zu wünschen übrig lässt, werden wir über die nächsten Jahre einen Shuttle-Service an den An- und Abreisetagen etablieren.

Auch für Hilfe in der aktuellen Krisenphase steht der Verein bereit.

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 under the strategic supervision of a foundation council (see Fig. 12.1). The association is chaired by a board (see Fig. 12.2 and Fig. 12.3).

The foundation assets are invested securely with the help of a professional asset management company specializing in foundations. The association has used the first years of its existence to build up a financial basis, and is now beginning to initiate the first initiatives to support Schloss Dagstuhl:

- Schloss Dagstuhl has also become known over the past decades for combining art and computer science. This has always been particularly inspiring for researchers visiting Dagstuhl. For this reason, the association has decided to support the art activities of Schloss Dagstuhl as permanently as possible. As a first step, the association has financed the purchase of two works by artists from the local region. In addition, funds are available for the professional inspection and archiving of the art collection and, if necessary, the preparation of a proposal for the hanging of the works.
- Another contribution of the association was made possible by the generous donation from the estate of a deceased friend of Schloss Dagstuhl: Since the connection from Schloss Dagstuhl via public transport to the Türkismühle train station often leaves much to be desired by our seminar participants, the association will establish a shuttle service on arrival and departure days over the next few years.

The association is also available for help in the current crisis phase.

Further information about the association as well as the membership application form can be found at http://www.friends-of-dagstuhl.de.

Stiftungsrat | Foundation council
Prof. Dr. Holger Hermanns (Vorstandsvorsitzender des Vereins “Friends of Dagstuhl” | First deputy chairperson of the association “Friends of Dagstuhl”) Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Dr. h.c. mult. Kurt Mehlhorn Max Planck Institute for Informatics (MPII), Saarbrücken, Germany
Prof. Dr. Dorothea Wagner Karlsruher Institut für Technologie (KIT), Germany

Fig. 12.1
Der Stiftungsrat der Stiftung “Informatik-Zentrum Schloss Dagstuhl”
The council of the foundation “Informatik-Zentrum Schloss Dagstuhl”

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. 12.2
Der Vorstand des Vereins “Friends of Dagstuhl”
The chair of the association “Friends of Dagstuhl”

Fig. 12.3
Der Vorstand des Vereins “Friends of Dagstuhl”, v.l.n.r.: Prof. Dr. Holger Hermanns, Angelika Müller-von Brochowski, und Erich Reindel.
The chair of the association “Friends of Dagstuhl”, f.l.t.r.: Prof. Holger Hermanns, Angelika Müller-von Brochowski, and Erich Reindel.

13 Statistiken
Statistics

Statistiken zu Seminaren und Workshops

13.1

Statistics on Seminars and Workshops

In diesem Abschnitt werden statistische Daten zum wissenschaftlichen Programm und der Zusammenstellung der Teilnehmer aufgeführt. Die Diagramme und Tabellen sind dabei wie nachfolgend beschrieben gegliedert.

Antrags-bezogene Daten: Die Anzahl eingereichter Anträge von Dagstuhl Seminaren und Dagstuhl Perspektiven Workshops sowie deren Akzeptanzraten sind in Fig. 13.1 dargestellt. Fig. 13.2 zeigt, wie die akzeptierten Seminare und Workshops sich bezüglich Größe und Länge aufgliedern.

Veranstaltungs-bezogene Daten: Fig. 13.3 zeigt Anzahl und Anteil der eingeladenen Seminarteilnehmer, welche die Einladung annehmen bzw. ablehnen. Die Verteilung dieser Annahmerate ist in Fig. 13.4 dargestellt. Fig. 13.5 zeigt dagegen, wie viel Prozent der zugesagten Größe (gemessen an der Personenanzahl) tatsächlich von einem Seminar belegt wurde. Daten zu Anzahl, Größe und Dauer der durchgeführten Seminare sind in Fig. 13.6 angegeben. Fig. 13.7 zeigt die Anzahl der verschiedenen Veranstaltungstypen.

Teilnehmer-bezogene Daten: Die Teilnehmerzahlen – abhängig vom Veranstaltungstyp – gibt Fig. 13.8 an. Fig. 13.9 zeigt die Verteilung der Herkunftsländer unserer Gäste.

Umfrage-bezogene Daten: Hier stellen wir ausgewählte Daten unserer fortlaufenden Befragung von Teilnehmern an Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops dar. Ein Überblick über die Ergebnisse der regelmäßigen Gästebefragungen kann Fig. 13.10 entnommen werden. Die Anzahl von früheren Seminarbesuchen kann man Fig. 13.11 entnehmen. Fig. 13.12 gibt Auskunft über die Altersstruktur der Teilnehmer. Während Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops sich primär an Forscher aus Universitäten und Forschungseinrichtungen richten, sind auch Anwender und Forscher aus der Industrie stets willkommen. Die Verteilung ihres Anteils ist in Fig. 13.13 gezeigt.

Auslastungs-bezogene Daten: Die Auslastung des Zentrums wird schließlich in Fig. 13.14 an Hand der Übernachtungen und ihrer Verteilung über die einzelnen Wochen getrennt nach Veranstaltungstypen aufgezeigt.

Geschlechter-bezogene Daten: Fig. 13.15 enthält Daten zur Geschlechter-Verteilung in der Seminarleitung. Dagegen zeigt Fig. 13.16 die Quote von Frauen bei der Beantragung von Seminaren sowohl bezüglich der Teams als auch bezüglich der gesamten Antragsteller. Die Abbildungen Fig. 13.17 und Fig. 13.18 zeigen insbesondere die Anteile weiblicher Teilnehmer bzw. Einladungen an weibliche Wissenschaftler. Die Verteilung der Rate der weiblichen Teilnehmer in den einzelnen Seminaren wird in Fig. 13.19 aufgezeigt.

This section provides statistical data about the scientific program and the composition of program participants. Charts and tables in this chapter may be outlined as follows.

Proposal-related data: Fig. 13.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, as well as acceptance rates for recent years. The size and duration of accepted seminars and workshops are displayed in Fig. 13.2.

Event-related data: Fig. 13.3 shows the number and the fraction of invited seminar participants who accepted or declined the invitation. The distribution of the rate is given in Fig. 13.4. In contrast, Fig. 13.5 visualizes how much of the reserved space was actually used by seminar participants. Data related to the number of seminars held in the last years together with their sizes and durations are given in Fig. 13.6. Fig. 13.7 shows the distribution of different types of events at Dagstuhl.

Participant-related data: Fig. 13.8 shows the number of participants according to event type. Fig. 13.9 shows the distribution of country affiliations.

Survey-related data: In this section we present data obtained from our ongoing Dagstuhl Seminar and Dagstuhl Perspectives Workshop guest survey project. An overview of the results of the 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.

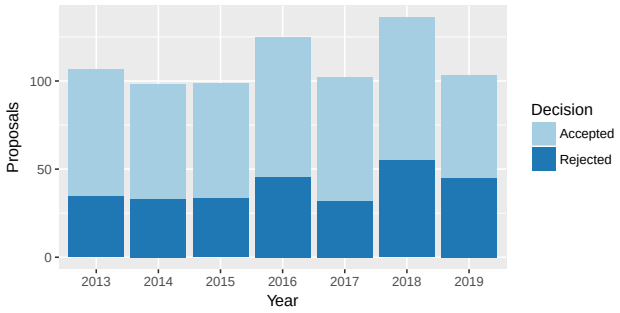

(a) Chart for 2013–2019

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
2013	107	72	67.3	35	32.7
2014	98	65	66.3	33	33.7
2015	99	65	65.7	34	34.3
2016	125	79	63.2	46	36.8
2017	102	70	68.6	32	31.4
2018	136	81	59.6	55	40.4
2019	103	58	56.3	45	43.7

(b) Detailed numbers for 2013–2019

Fig. 13.1
Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

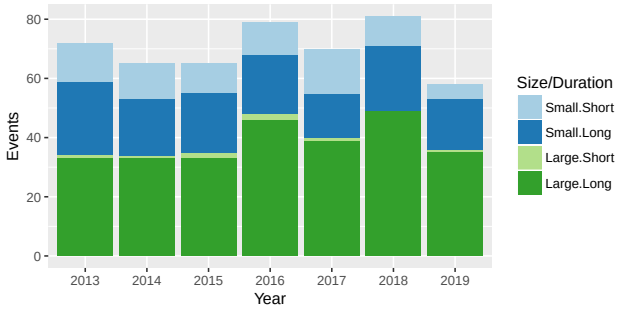

(a) Chart for 2013–2019

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2013	13	25	1	33	72
2014	12	19	1	33	65
2015	10	20	2	33	65
2016	11	20	2	46	79
2017	15	15	1	39	70
2018	10	22	0	49	81
2019	5	17	1	35	58

(b) Detailed numbers for 2013–2019

Fig. 13.2
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2013–2019. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

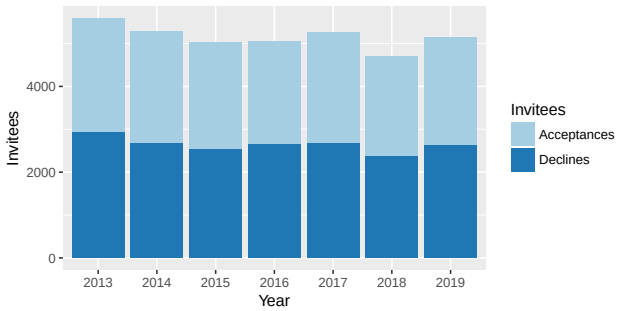

(a) Chart for 2013–2019

Year	Invitees	Acceptances		Declines	
	#	#	%	#	%
2013	5591	2639	47.2	2952	52.8
2014	5285	2590	49.0	2695	51.0
2015	5023	2473	49.2	2550	50.8
2016	5060	2393	47.3	2667	52.7
2017	5267	2572	48.8	2695	51.2
2018	4692	2320	49.4	2372	50.6
2019	5143	2498	48.6	2645	51.4

(b) Detailed numbers for 2013–2019

Fig. 13.3
Total number of invitees, acceptances, and declines for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

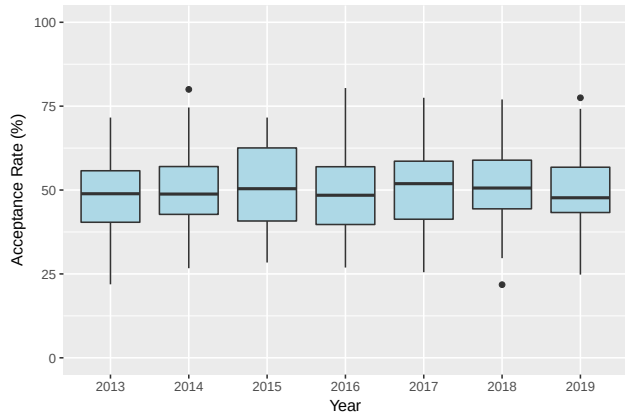

(a) Chart for 2013–2019

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2013	21.9	71.6	48.4	11.2
2014	26.7	80.0	50.2	11.2
2015	28.4	71.6	50.7	12.4
2016	26.9	80.4	48.6	11.2
2017	25.5	77.5	50.3	12.4
2018	21.8	77.0	51.2	12.0
2019	24.8	77.5	49.8	11.4

(b) Detailed numbers for 2013–2019

Fig. 13.4
Distribution of the acceptance rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2013–2019. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

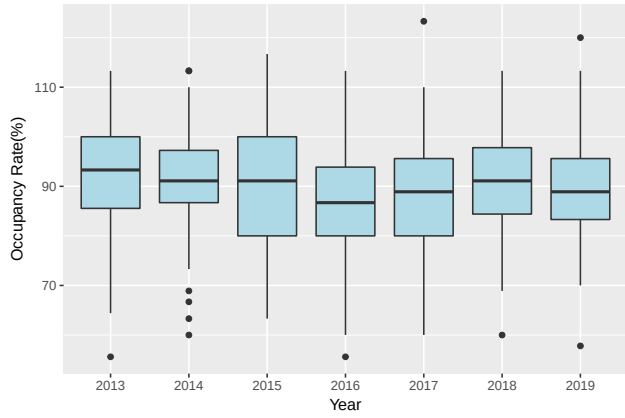

(a) Chart for 2013–2019

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2013	55.6	113.3	92.1	12.2
2014	60.0	113.3	90.6	10.3
2015	63.3	116.7	89.6	12.5
2016	55.6	113.3	86.7	11.8
2017	60.0	123.3	87.3	12.3
2018	60.0	113.3	90.3	10.2
2019	57.8	120.0	89.1	10.7

(b) Detailed numbers for 2013–2019

Fig. 13.5
Distribution of the occupancy rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2013–2019. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

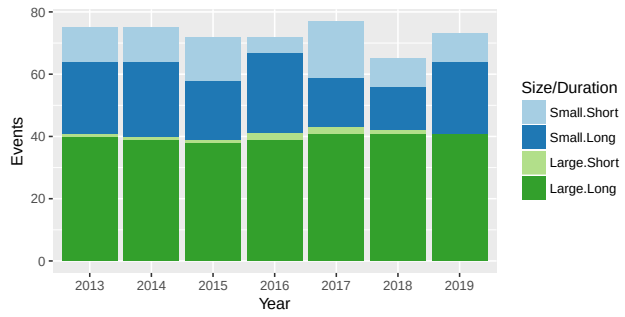

(a) Chart for 2013–2019

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2013	11	23	1	40	75
2014	11	24	1	39	75
2015	14	19	1	38	72
2016	5	26	2	39	72
2017	18	16	2	41	77
2018	9	14	1	41	65
2019	9	23	0	41	73

(b) Detailed numbers for 2013–2019

Fig. 13.6
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2013–2019. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

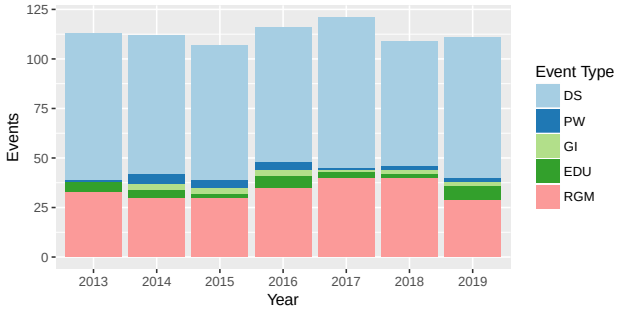

(a) Chart for 2013–2019

Year	DS	PW	GI	EDU	RGM	Total
2013	74	1	0	5	33	113
2014	70	5	3	4	30	112
2015	68	4	3	2	30	107
2016	68	4	3	6	35	116
2017	76	1	1	3	40	121
2018	63	2	2	2	40	109
2019	71	2	2	7	29	111

(b) Detailed numbers for 2013–2019

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.

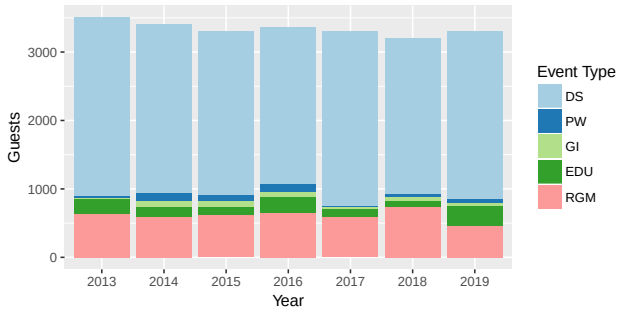

(a) Chart for 2013–2019

Year	DS		PW		GI		EDU		RGM		Total
	#	%	#	%	#	%	#	%	#	%	#
2013	2610	74.5	29	0.8	0	0.0	230	6.6	634	18.1	3503
2014	2463	72.2	127	3.7	86	2.5	144	4.2	589	17.3	3409
2015	2385	72.3	88	2.7	90	2.7	111	3.4	624	18.9	3298
2016	2280	68.0	113	3.4	78	2.3	232	6.9	650	19.4	3353
2017	2551	77.1	21	0.6	21	0.6	131	4.0	584	17.7	3308
2018	2268	70.8	52	1.6	50	1.6	99	3.1	733	22.9	3202
2019	2450	74.3	48	1.5	50	1.5	282	8.5	469	14.2	3299

(b) Detailed numbers for 2013–2019

Fig. 13.8
Number of participants. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

Country	A	B	Total
Germany	535	559	1094
United States	523	39	562
United Kingdom	238	32	270
France	173	13	186
Netherlands	111	25	136
Austria	85	13	98
Canada	89	8	97
Italy	72	8	80
Israel	68	2	70
Sweden	62	5	67
Switzerland	52	14	66
Belgium	45	2	47
Spain	39	8	47
Japan	40	5	45
Luxembourg	17	27	44
Denmark	39	2	41
Australia	35	4	39
Finland	24	2	26
India	17	6	23
Portugal	20	3	23
Czech Republic	18	1	19
Norway	17	2	19
Singapore	17	2	19
China	15	3	18
Poland	17	1	18
Brazil	12	0	12
Ireland	8	4	12
Republic of Korea	9	2	11
Chile	10	0	10
Hungary	9	1	10
New Zealand	10	0	10
Taiwan	10	0	10
Russian Federation	8	0	8
Hong Kong	6	1	7
Slovenia	5	0	5
Estonia	4	0	4
Iceland	4	0	4
Serbia	0	4	4
Turkey	4	0	4
United Arab Emirates	4	0	4
Mexico	3	0	3
Argentina	2	0	2
Croatia	0	2	2
Cyprus	2	0	2
Egypt	1	1	2
Greece	2	0	2
Romania	2	0	2
Rwanda	2	0	2
Slovak Republic	2	0	2
South Africa	2	0	2
Uganda	2	0	2
Kazakhstan	1	0	1
Latvia	1	0	1
Malta	1	0	1
Morocco	1	0	1
Saudi Arabia	1	0	1
Somalia	1	0	1
Tunisia	1	0	1
Total	2498	801	3299

(a) Details for 2019 by country

Region	A		B		Total	
	#	%	#	%	#	%
Europe (w/o Germany)	1081	43.3	169	21.1	1250	37.9
Germany	535	21.4	559	69.8	1094	33.2
North America	612	24.5	47	5.9	659	20
Asia	188	7.5	21	2.6	209	6.3
Australia	45	1.8	4	0.5	49	1.5
South America	27	1.1	0	0	27	0.8
Africa	10	0.4	1	0.1	11	0.3
Total	2498	100	801	100	3299	100

(b) Details for 2019 by region

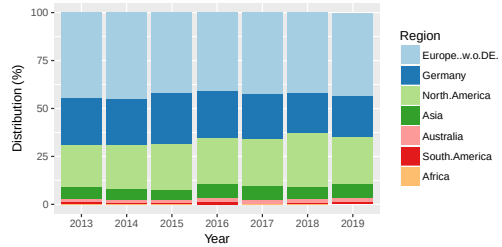

(c) Graphical distribution of seminar type A in 2013–2019 by year and region

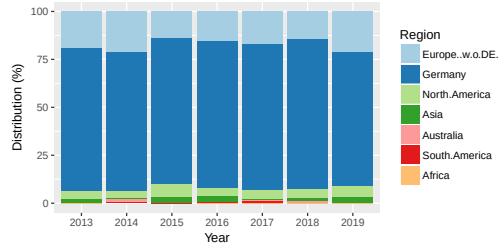

(d) Graphical distribution of seminar type B in 2013–2019 by year and region

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).

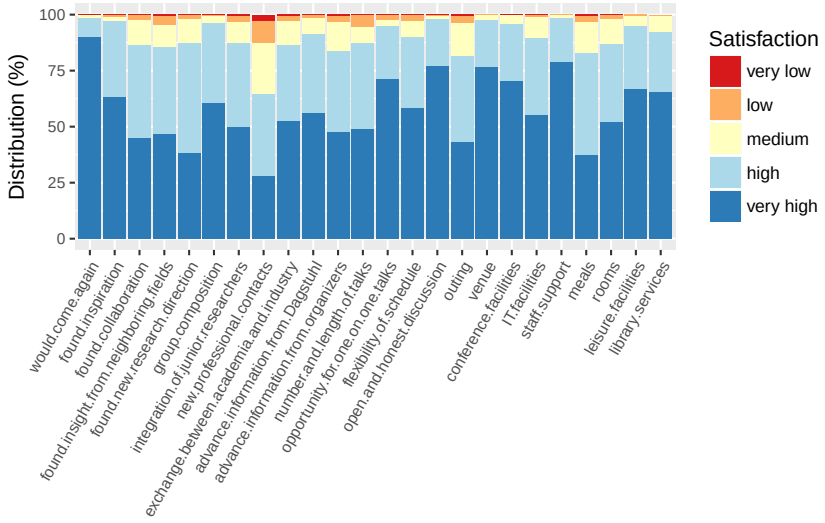

(a) Graphical distribution for 2019

	2013	2014	2015	2016	2017	2018	2019	2019 – 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	3	17	127	1361	1509
found inspiration	4.5	4.4	4.5	4.5	4.5	4.6	4.6	4	9	30	513	964	1520
found collaboration	4.0	4.1	4.1	4.1	4.2	4.2	4.3	3	28	169	628	675	1503
found insight from neighboring fields	4.1	4.2	4.3	4.2	4.2	4.2	4.3	9	55	153	587	710	1514
found new research direction	4.0	4.0	4.1	4.1	4.1	4.2	4.2	2	26	157	740	575	1500
group composition	4.4	4.4	4.5	4.5	4.5	4.5	4.6	1	9	47	537	922	1516
integration of junior researchers	4.2	4.2	4.2	4.3	4.3	4.3	4.3	7	39	141	567	747	1501
new professional contacts	3.7	3.7	3.6	3.7	3.8	3.8	3.8	38	148	336	540	414	1476
exchange between academia and industry	4.2	4.2	4.3	4.3	4.4	4.4	4.4	6	19	107	337	523	992
advance information from Dagstuhl	4.4	4.4	4.4	4.4	4.4	4.4	4.5	2	15	111	519	839	1486
advance information from organizers	4.1	4.1	4.1	4.2	4.1	4.3	4.3	9	38	190	528	702	1467
number and length of talks	4.2	4.1	4.2	4.3	4.2	4.3	4.3	3	74	112	577	741	1507
opportunity for one on one talks	4.5	4.5	4.5	4.6	4.6	4.6	4.6	1	33	44	350	1078	1506
flexibility of schedule	4.2	4.3	4.3	4.4	4.3	4.3	4.5	4	34	110	470	877	1495
open and honest discussion	4.7	4.7	4.7	4.7	4.7	4.7	4.7	2	6	22	313	1155	1498
outing	4.1	4.1	4.1	4.2	4.2	4.3	4.2	8	33	171	448	502	1162
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.7	0	2	29	318	1153	1502
conference facilities	4.6	4.7	4.6	4.7	4.7	4.7	4.7	1	2	58	383	1058	1502
IT facilities	4.4	4.4	4.3	4.4	4.3	4.4	4.4	1	9	126	460	736	1332
staff support	4.7	4.7	4.7	4.7	4.7	4.8	4.8	0	1	18	287	1150	1456
meals	4.1	4.1	4.1	4.1	4.1	4.1	4.2	8	41	208	681	563	1501
rooms	4.4	4.4	4.4	4.4	4.4	4.4	4.4	2	22	168	526	782	1500
leisure facilities	4.6	4.6	4.6	4.5	4.5	4.6	4.6	0	10	57	388	925	1380
library services	4.5	4.5	4.5	4.5	4.5	4.6	4.6	2	1	52	185	456	696

(b) Averages for 2013–2019 and detailed numbers for 2019: 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.

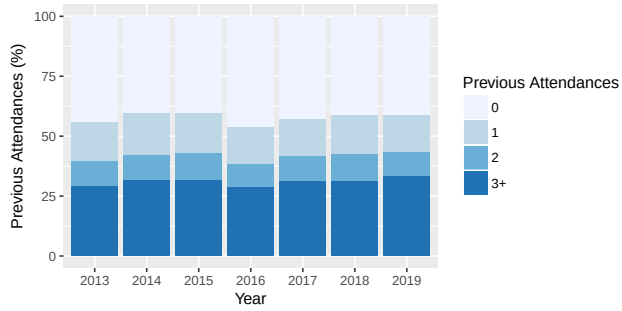

(a) Graphical distribution for 2013–2019

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	#
2013	630	44	237	17	145	10	422	29	1434
2014	561	40	239	17	144	10	443	32	1387
2015	573	40	234	17	158	11	451	32	1416
2016	654	46	217	15	137	10	410	29	1418
2017	607	43	222	16	148	10	446	31	1423
2018	557	41	219	16	148	11	425	32	1349
2019	615	41	230	15	144	10	503	34	1492

(b) Detailed numbers for 2013–2019

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.

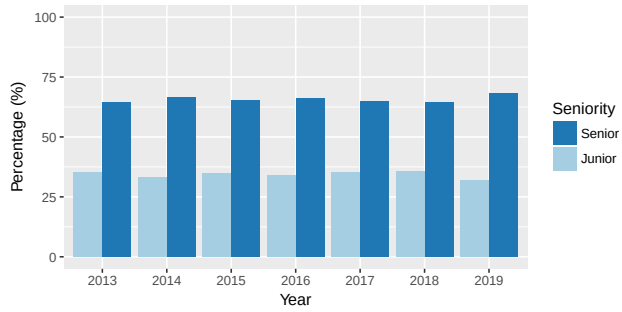

(a) Chart for 2013–2019

Year	Junior		Senior		Total
	#	%	#	%	#
2013	413	35.4	754	64.6	1167
2014	382	33.3	765	66.7	1147
2015	410	34.9	764	65.1	1174
2016	404	33.9	787	66.1	1191
2017	422	35.2	778	64.8	1200
2018	401	35.7	722	64.3	1123
2019	385	31.9	823	68.1	1208

(b) Detailed numbers for 2013–2019

Fig. 13.12
Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

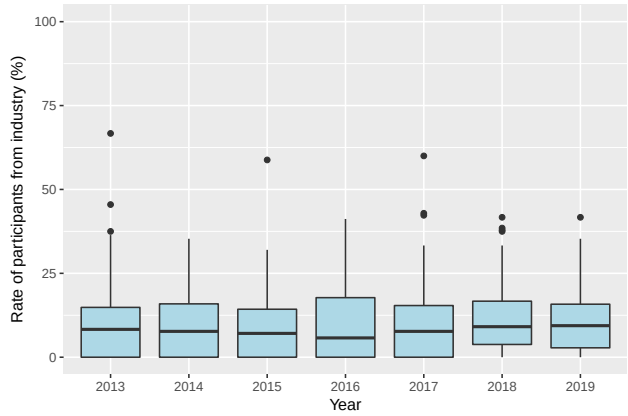

(a) Chart for 2013–2019

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2013	0.0	66.7	11.6	12.8
2014	0.0	35.3	9.4	9.4
2015	0.0	58.8	9.8	10.5
2016	0.0	41.2	10.3	11.0
2017	0.0	60.0	10.9	11.6
2018	0.0	41.7	11.1	10.4
2019	0.0	41.7	11.4	10.7

(b) Detailed numbers for 2013–2019

Fig. 13.13
Distribution of the rate of participants with self-assigned primary occupation in business per Dagstuhl Seminar and Dagstuhl Perspectives Workshop in 2013–2019, 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".

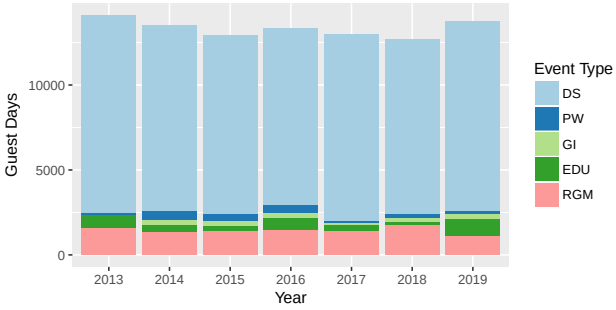

(a) Chart for 2013–2019

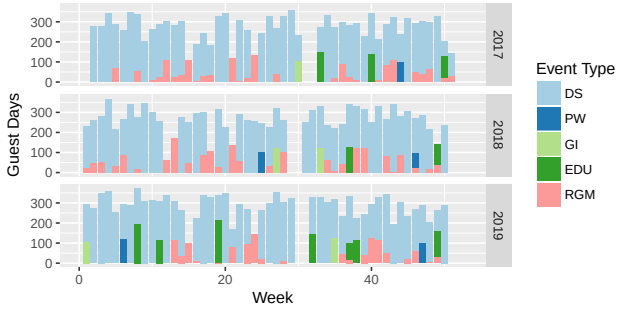

(b) Graphical distribution for 2017–2019 by week

Year	DS	PW	GI	EDU	RGM	Total
2013	11612	130	0	753	1614	14109
2014	10939	475	348	390	1370	13522
2015	10491	380	344	261	1424	12900
2016	10362	495	315	703	1462	13337
2017	10989	102	105	401	1391	12988
2018	10270	182	250	231	1740	12673
2019	11127	225	239	1004	1144	13739

(c) Detailed numbers for 2013–2019

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.

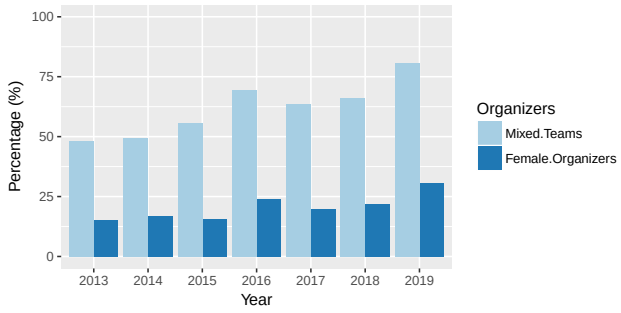

(a) Chart for 2013–2019

	Organizer Teams			Organizers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
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
2018	65	43	66.2	259	57	22.0
2019	73	59	80.8	291	89	30.6

(b) Detailed numbers for 2013–2019

Fig. 13.15

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

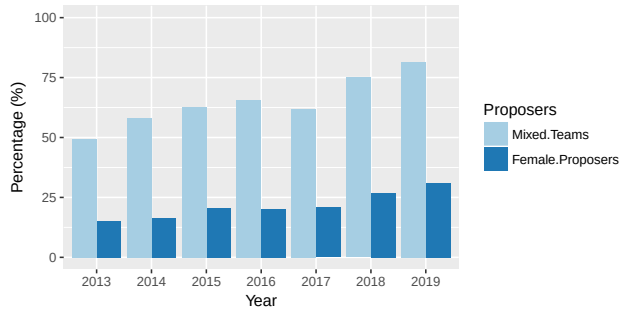

(a) Chart for 2013–2019

	Proposer Teams			Proposers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2013	107	53	49.5	431	66	15.3
2014	98	57	58.2	387	64	16.5
2015	99	62	62.6	391	80	20.5
2016	125	82	65.6	491	99	20.2
2017	102	63	61.8	394	82	20.8
2018	136	102	75.0	522	140	26.8
2019	103	84	81.6	411	127	30.9

(b) Detailed numbers for 2013–2019

Fig. 13.16
Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with mixed-gender proposer teams.

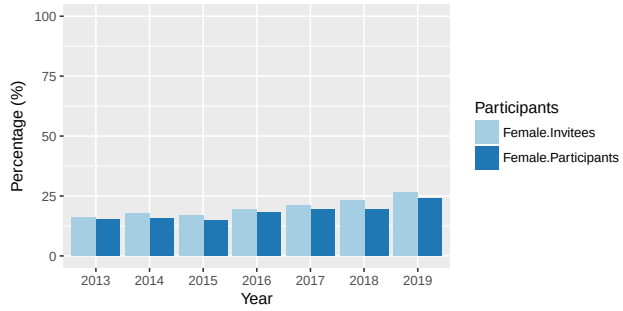

(a) Chart for 2013–2019

	Invitees			Participants		
Year	Total	Female		Total	Female	
	#	#	%	#	#	%
2013	5591	890	15.9	2639	402	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	437	18.3
2017	5267	1114	21.2	2572	497	19.3
2018	4692	1089	23.2	2320	455	19.6
2019	5143	1369	26.6	2498	603	24.1

(b) Detailed numbers for 2013–2019

Fig. 13.17
Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.

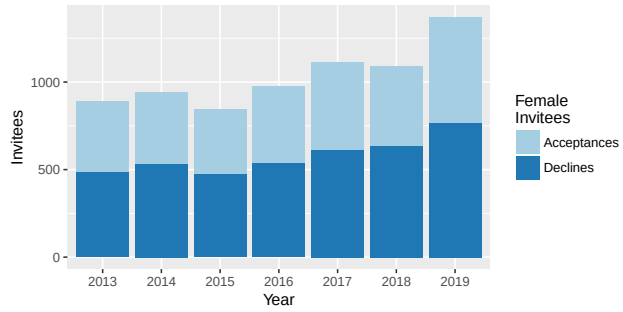

(a) Chart for 2013–2019

Year	Female Invitees	Acceptances		Declines	
	#	#	%	#	%
2013	890	402	45.2	488	54.8
2014	943	406	43.1	537	56.9
2015	845	369	43.7	476	56.3
2016	977	437	44.7	540	55.3
2017	1114	497	44.6	617	55.4
2018	1089	455	41.8	634	58.2
2019	1369	603	44.0	766	56.0

(b) Detailed numbers for 2013–2019

Fig. 13.18
Female invitees to Dagstuhl Seminar and Dagstuhl Perspectives Workshops.

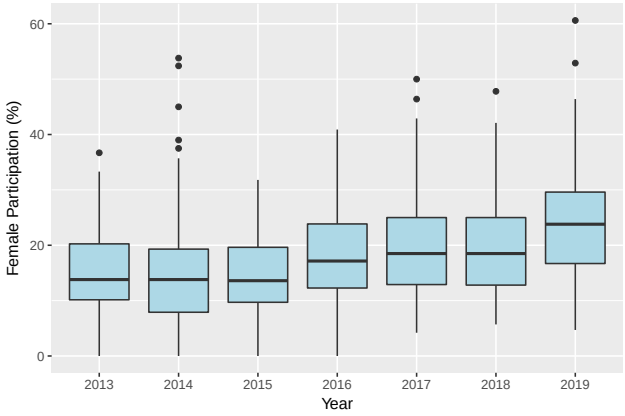

(a) Chart for 2013–2019

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2013	0.0	36.7	15.2	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	19.8	9.8
2018	5.7	47.8	20.0	9.3
2019	4.7	60.6	24.6	10.6

(b) Detailed numbers for 2013–2019

Fig. 13.19
Distribution of female participants rate per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2013–2019. 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.

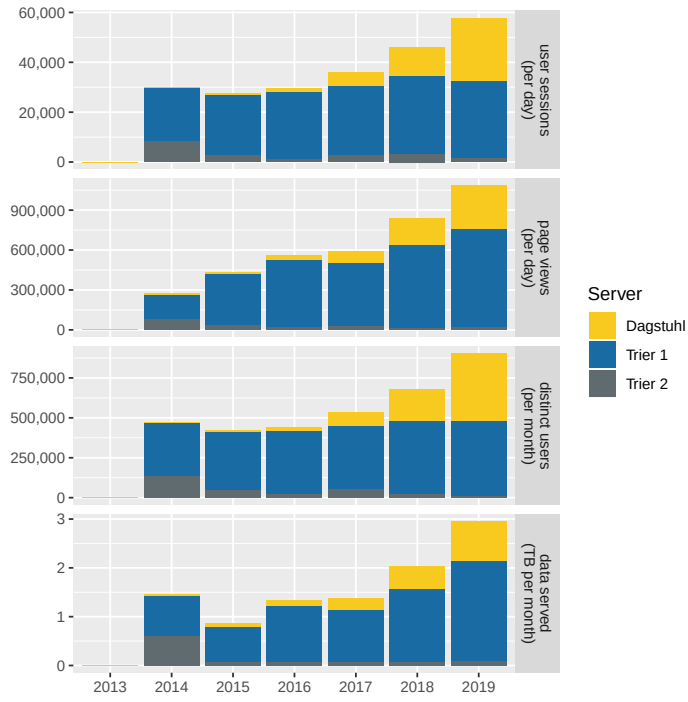

(a) Chart for 2014–2019

	Trier 1		Trier 2		Dagstuhl		Total		
	2018	2019	2018	2019	2018	2019	2018	2019	%
user sessions (visits) per day	31,530	31,024	3,233	1,808	11,483	24,994	46,247	57,827	+25.0
page views per day	618,067	735,190	20,208	22,761	202,301	326,053	840,577	1,084,005	+29.0
page views per user session	19.6	23.7	6.2	12.6	17.6	13.0	18.2	18.7	+3.1
distinct users (IPs) per month	451,769	466,015	27,448	12,963	197,270	424,106	676,489	903,085	+33.5
data served per month	1,535.0 GB	2,114.1 GB	72.6 GB	89.6 GB	469.7 GB	821.3 GB	2,077.3 GB	3,025.0 GB	+45.6

(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.

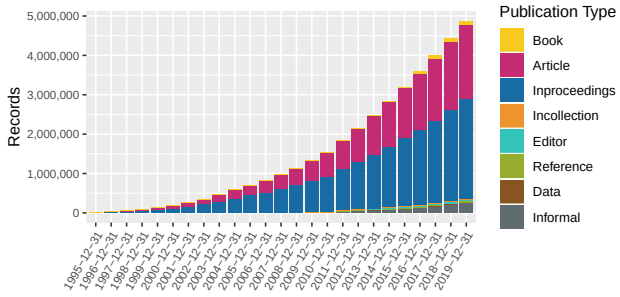

(a) Chart for 1996–2019

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2013	16,819	0.7	997,820	40.2	1,350,713	54.4	12,797	0.5	22,771	0.9	13,125	0.5	0	0.0	69,905	2.8	2,483,950
2014	17,533	0.6	1,129,231	39.8	1,545,065	54.5	14,470	0.5	26,137	0.9	14,690	0.5	0	0.0	88,217	3.1	2,835,343
2015	18,318	0.6	1,281,245	40.0	1,724,262	53.9	16,288	0.5	30,044	0.9	19,103	0.6	12	0.0	110,974	3.5	3,200,246
2016	51,070	1.4	1,429,427	39.7	1,912,895	53.1	19,774	0.5	33,782	0.9	20,174	0.6	26	0.0	134,354	3.7	3,601,502
2017	77,408	1.9	1,576,972	39.4	2,091,486	52.2	23,101	0.6	37,049	0.9	23,089	0.6	49	0.0	174,723	4.4	4,003,877
2018	83,249	1.9	1,725,704	39.0	2,306,585	52.2	24,708	0.6	40,795	0.9	23,150	0.5	514	0.0	216,984	4.9	4,421,689
2019	90,826	1.9	1,896,014	38.8	2,518,298	51.6	30,457	0.6	44,898	0.9	26,997	0.6	1,402	0.0	271,633	5.6	4,880,525

(b) Detailed numbers for 2013–2019

Fig. 13.21
Development of the total size of the dblp database.

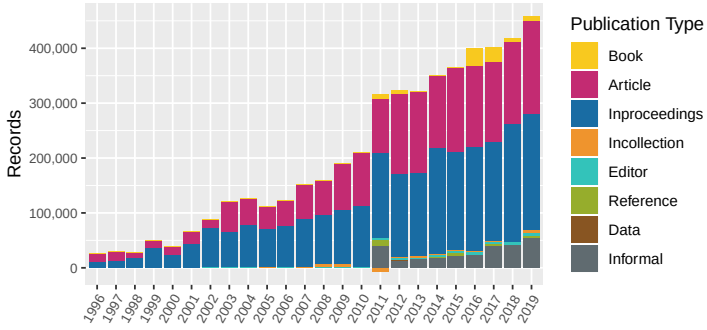

(a) Chart for 1996–2019

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2013	782	0.2	147,217	45.9	150,821	47.0	3,166	1.0	2,617	0.8	0	0.0	0	0.0	16,095	5.0	320,698
2014	714	0.2	131,411	37.4	194,352	55.3	1,673	0.5	3,366	1.0	1,565	0.4	0	0.0	18,312	5.2	351,393
2015	785	0.2	152,014	41.7	179,197	49.1	1,818	0.5	3,907	1.1	4,413	1.2	12	0.0	22,757	6.2	364,903
2016	32,752	8.2	148,182	36.9	188,633	47.0	3,486	0.9	3,738	0.9	1,071	0.3	14	0.0	23,380	5.8	401,256
2017	26,338	6.5	147,545	36.7	178,591	44.4	3,327	0.8	3,267	0.8	2,915	0.7	23	0.0	40,369	10.0	402,375
2018	5,841	1.4	148,732	35.6	215,099	51.5	1,607	0.4	3,746	0.9	61	0.0	465	0.1	42,261	10.1	417,812
2019	7,577	1.7	170,310	37.1	211,713	46.1	5,749	1.3	4,103	0.9	3,847	0.8	888	0.2	54,649	11.9	458,836

(b) Detailed numbers for 2013–2019

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*.

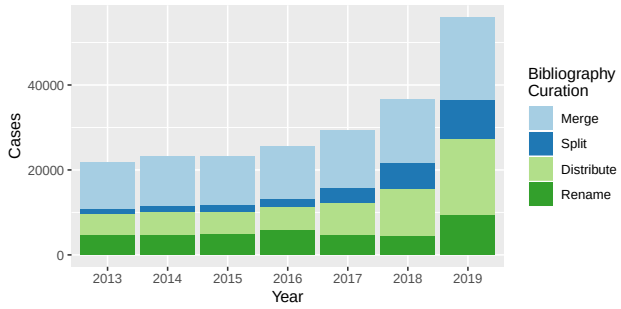

(a) Chart for 2013–2019

Year	Merge		Split		Distribute		Rename		Total
	#	%	#	%	#	%	#	%	#
2013	10,863	49.8	1,326	6.1	4,891	22.4	4,744	21.7	21,824
2014	11,564	50.0	1,382	6.0	5,410	23.4	4,768	20.6	23,124
2015	11,526	49.6	1,495	6.4	5,323	22.9	4,876	21.0	23,220
2016	12,426	48.4	1,913	7.5	5,310	20.7	5,999	23.4	25,648
2017	13,537	46.0	3,660	12.4	7,465	25.3	4,786	16.3	29,448
2018	14,906	40.6	6,282	17.1	11,014	30.0	4,524	12.3	36,726
2019	19,454	34.8	9,194	16.4	17,766	31.7	9,555	17.1	55,969

(b) Detailed numbers for 2013–2019

Fig. 13.23

Curation of existing dblp author bibliographies. The figures give the number of distinct edit cases (measured between the first and the last day of every given year) where a dblp team member manually corrected the assignment of publications within dblp author bibliographies. We distinguish between four curation cases: *Merge* = Two or more synonymous bibliographies have been merged into a single bibliography. *Split* = A single, homonymous bibliography has been split into two or more bibliographies. *Distribute* = A mixed case where records from two or more bibliographies have been redistributed between two or more bibliographies. *Rename* = A case where no actual publications have been reassigned, but the surface form of the author name(s) of a bibliography have been corrected or improved. These figures correct flawed figures given in earlier reports.

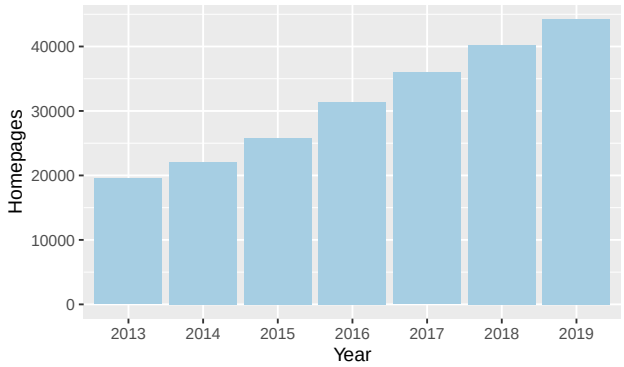

(a) Chart for 2013–2019

Year	Homepages
2013	19,511
2014	22,101
2015	25,814
2016	31,353
2017	35,973
2018	40,179
2019	44,248

(b) Detailed numbers for 2013–2019

Fig. 13.24

Linked and verified academic homepages in dblp author bibliographies. A single author bibliography may be linked to multiple academic homepages. These figures exclude linked external IDs which are given in Figure 13.25.

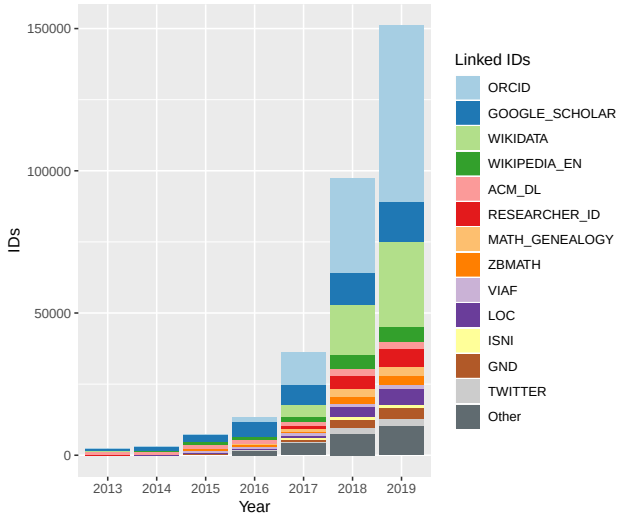

(a) Chart for 2013–2019

Year	ORCID		Google Scholar		Wikidata		Wikipedia (en)		ACM DL		ResearcherID		Math Genealogy	
	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2013	14	0.6	947	39.9	0	0.0	207	8.7	1,169	49.2	2	0.1	0	0.0
2014	24	0.8	1,379	46.9	0	0.0	271	9.2	1,173	39.9	8	0.3	0	0.0
2015	89	1.2	2,510	34.6	0	0.0	1,002	13.8	1,225	16.9	34	0.5	341	4.7
2016	1,717	12.8	4,999	37.3	4	0.0	1,234	9.2	1,236	9.2	154	1.2	420	3.1
2017	11,591	31.8	7,326	20.1	4,046	11.1	1,550	4.3	1,425	3.9	1,297	3.6	879	2.4
2018	33,185	34.1	11,226	11.5	17,405	17.9	5,065	5.2	2,339	2.4	4,824	5.0	2,717	2.8
2019	61,976	41.1	13,726	9.1	30,022	19.9	5,547	3.7	2,448	1.6	6,192	4.1	3,071	2.0

Year	Zentralblatt MATH		VIAF		LOC		ISNI		GND		Twitter		Other		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2013	0	0.0	2	0.1	0	0.0	0	0.0	0	0.0	1	0.0	32	1.3	2,374
2014	4	0.1	3	0.1	0	0.0	1	0.0	0	0.0	13	0.4	64	2.2	2,940
2015	622	8.6	573	7.9	357	4.9	1	0.0	240	3.3	67	0.9	200	2.8	7,261
2016	687	5.1	585	4.4	357	2.7	7	0.1	243	1.8	184	1.4	1,562	11.7	13,389
2017	701	1.9	945	2.6	677	1.9	671	1.8	481	1.3	453	1.2	4,419	12.1	36,461
2018	2,435	2.5	1,057	1.1	3,320	3.4	1,381	1.4	2,793	2.9	1,764	1.8	7,773	8.0	97,284
2019	3,363	2.2	1,060	0.7	5,680	3.8	1,383	0.9	3,782	2.5	2,183	1.4	10,532	7.0	150,965

(b) Detailed numbers for 2013–2019

Fig. 13.25
Linked and verified external person IDs in dblp author bibliographies. A single bibliography 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.

Die Kennzahlen der Zeitschrift LITES können Fig. 13.31 entnommen werden.

Die verschiedenen Publikationsserien wurden in unterschiedlichen Jahren zwischen 2009 und 2015 gegründet. Wir stellen in den Statistiken dennoch stets den gesamten Zeitraum (2013–2019) 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 of the service-focused series are presented afterwards. Fig. 13.29 presents numbers related to OASIcs and Fig. 13.30 numbers related to LIPIcs.

We summarize the publications of the journal LITES in Fig. 13.31.

Please note that the publication series were established in different years in the period between 2009 and 2015. However, we always consider this complete period (2013–2019).

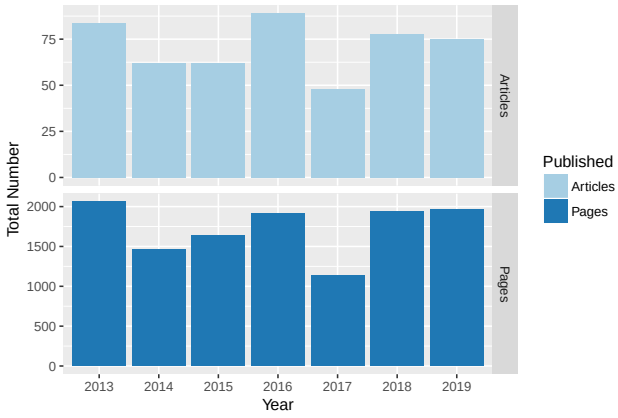

(a) Graphical distribution for 2013–2019

Year	Articles	Pages
2013	84	2059
2014	62	1464
2015	62	1636
2016	89	1910
2017	48	1138
2018	78	1938
2019	75	1961

(b) Detailed numbers for 2013–2019

Fig. 13.26
Statistics about Dagstuhl Reports published between 2013 to 2019.

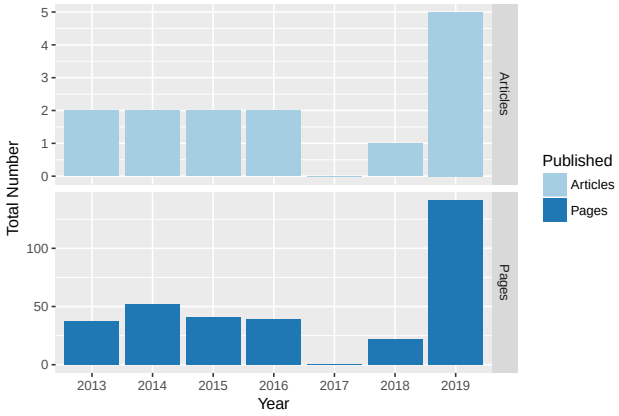

(a) Graphical distribution for 2013–2019

Year	Articles	Pages
2013	2	37
2014	2	52
2015	2	41
2016	2	39
2017	0	0
2018	1	22
2019	5	141

(b) Detailed numbers for 2013–2019

Fig. 13.27
Statistics about Dagstuhl Manifestos published between 2013 to 2019.

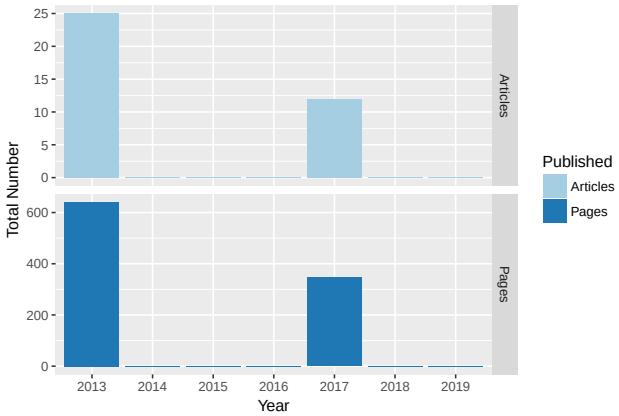

(a) Graphical distribution for 2013–2019

Year	Volumes	Articles	Pages
2013	3	25	641
2014	0	0	0
2015	0	0	0
2016	0	0	0
2017	1	12	346
2018	0	0	0
2019	0	0	0

(b) Detailed numbers for 2013–2019

Fig. 13.28
Statistics about Dagstuhl Follow-Ups volumes published between 2013 to 2019.

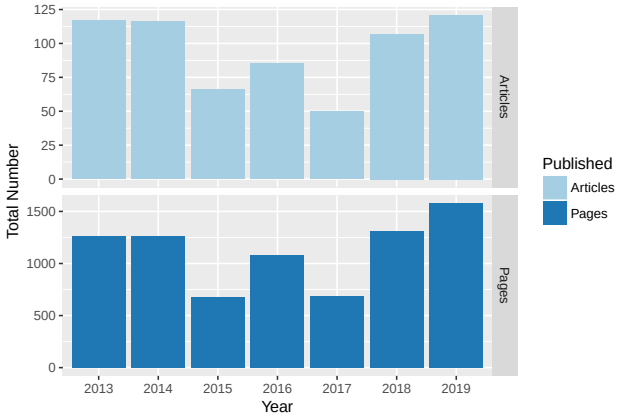

(a) Graphical distribution for 2013–2019

Year	Volumes	Articles	Pages
2013	7	117	1265
2014	8	116	1264
2015	6	66	674
2016	6	85	1078
2017	3	50	684
2018	7	107	1312
2019	9	121	1576

(b) Detailed numbers for 2013–2019

Fig. 13.29
Statistics about OASIcs volumes published between 2013 to 2019.

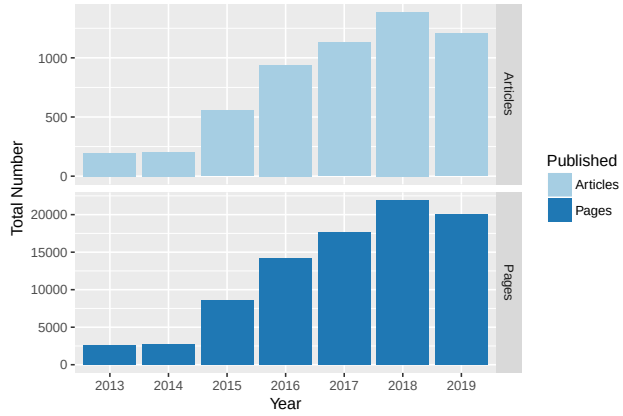

(a) Graphical distribution for 2013–2019

Year	Volumes	Articles	Pages
2013	6	195	2607
2014	5	204	2752
2015	16	553	8565
2016	19	939	14222
2017	25	1127	17687
2018	32	1387	21871
2019	29	1208	20032

(b) Detailed numbers for 2013–2019

Fig. 13.30
Statistics about LIPIcs volumes published between 2013 to 2019.

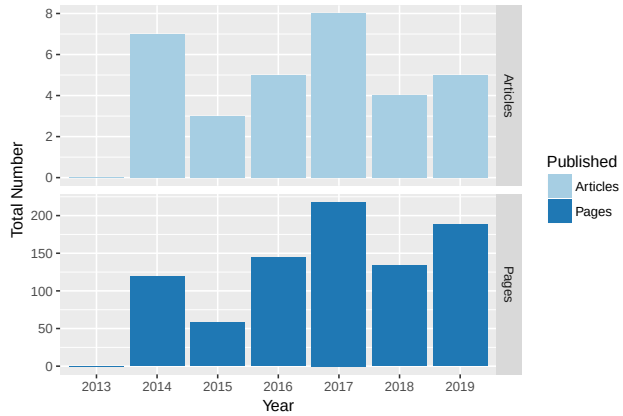

(a) Graphical distribution for 2013–2019

Year	Articles	Pages
2013	0	0
2014	7	119
2015	3	58
2016	5	144
2017	8	218
2018	4	134
2019	5	188

(b) Detailed numbers for 2013–2019

Fig. 13.31
Statistics about LITES articles published between 2013 to 2019.

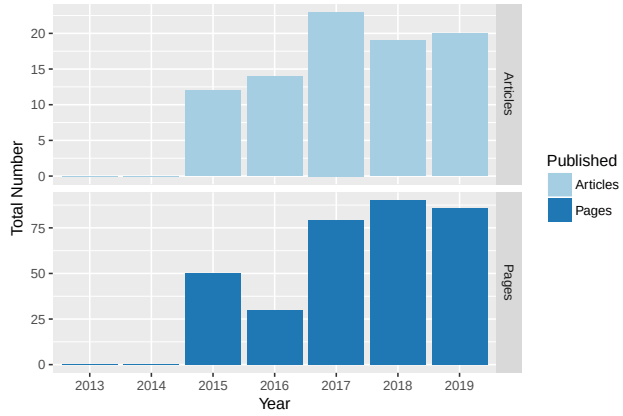

(a) Graphical distribution for 2013–2019

Year	Articles	Pages
2013	0	0
2014	0	0
2015	12	50
2016	14	30
2017	23	79
2018	19	90
2019	20	86

(b) Detailed numbers for 2013–2019

Fig. 13.32
Statistics about DARTS artifacts published between 2013 to 2019.

14 Veranstaltungen 2019
Schedule of Events 2019

Dagstuhl-Seminare 14.1 *Dagstuhl Seminars*

19021 – Joint Processing of Language and Visual Data for Better Automated Understanding
Yun Fu (Northeastern University – Boston, US), Marie-Francine Moens (KU Leuven, BE), Lucia Specia (Imperial College London, GB), Tinne Tuytelaars (KU Leuven, BE)
January 6–11, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19021

19031 – Logics for Dependence and Independence
Erich Grädel (RWTH Aachen, DE), Phokion G. Kolaitis (University of California – Santa Cruz, US), Juha Kontinen (University of Helsinki, FI), Heribert Vollmer (Leibniz Universität Hannover, DE)
January 13–18, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19031

19032 – Conditional Logics and Conditional Reasoning: New Joint Perspectives
Guillaume Aucher (University of Rennes 1 & IRISA Rennes, FR), Paul Egré (ENS – Paris, FR), Gabriele Kern-Isberner (TU Dortmund, DE), Francesca Poggiolesi (CNRS – Paris, FR)
January 13–16, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19032

19041 – New Horizons in Parameterized Complexity
Fedor V. Fomin (University of Bergen, NO), Dániel Marx (Hungarian Academy of Sciences – Budapest, HU), Saket Saurabh (Institute of Mathematical Sciences – Chennai, IN), Meirav Zehavi (Ben Gurion University – Beer Sheva, IL)
January 20–25, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19041

19042 – Practical Yet Composably Secure Cryptographic Protocols
Jan Camenisch (Dfinity Foundation – Zug, CH), Ralf Küsters (Universität Stuttgart, DE), Anna Lysyanskaya (Brown University – Providence, US), Alessandra Scafuro (North Carolina State University – Raleigh, US)
January 20–25, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19042

19051 – Data Structures for the Cloud and External Memory Data
Gerth Stølting Brodal (Aarhus University, DK), Ulrich Carsten Meyer (Goethe-Universität – Frankfurt am Main, DE), Markus E. Nebel (Universität Bielefeld, DE), Robert Sedgewick (Princeton University, US)
January 27 to February 1, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19051

19052 – Computational Methods for Melody and Voice Processing in Music Recordings
Emilia Gómez (UPF – Barcelona, ES), Meinard Müller (Universität Erlangen-Nürnberg, DE), Yi-Hsuan Yang (Academia Sinica – Taipei, TW)
January 27 to February 1, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19052

19061 – Visual Analytics of Multilayer Networks Across Disciplines
Nathalie Henry Riche (Microsoft Research – Redmond, US), Mikko Kivelä (Aalto University, FI), Fintan McGee (Luxembourg Inst. of Science & Technology, LU), Guy Melançon (University of Bordeaux, FR), Tatiana von Landesberger (TU Darmstadt, DE)
February 3–8, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19061

19062 – Bringing CP, SAT and SMT together: Next Challenges in Constraint Solving
Sébastien Bardin (CEA LIST, FR), Nikolaj S. Bjørner (Microsoft Research – Redmond, US), Cristian Cadar (Imperial College London, GB), Vijay Ganesh (University of Waterloo, CA)
February 3–6, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19062

19071 – Specification Formalisms for Modern Cyber-Physical Systems
Jyotirmoy Deshmukh (USC – Los Angeles, US), Oded Maler (VERIMAG – Grenoble, FR), Dejan Nickovic (AIT – Austrian Institute of Technology – Wien, AT)
February 10–15, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19071

19081 – Verification and Synthesis of Human-Robot Interaction
Rachid Alami (LAAS – Toulouse, FR), Kerstin I. Eder (University of Bristol, GB), Guy Hoffman (Cornell University, US), Hadas Kress-Gazit (Cornell University, US)
February 17–22, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19081

19082 – AI for the Social Good
Claudia Clopath (Imperial College London, GB), Ruben De Winne (Oxfam Novib – The Hague, NL), Mohammad Emtiyaz Khan (RIKEN – Tokyo, JP), Tom Schaul (Google DeepMind – London, GB)
February 17–22, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19082

19092 – Beyond-Planar Graphs: Combinatorics, Models and Algorithms
Seok-Hee Hong (The University of Sydney, AU), Michael Kaufmann (Universität Tübingen, DE), János Pach (EPFL – Lausanne, CH), Csaba D. Tóth (California State University – Northridge, US)
February 24 to March 1, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19092

19101 – Analysis, Design, and Control of Predictable Interconnected Systems
Kunal Agrawal (Washington University – St. Louis, US), Enrico Bini (University of Turin, IT), Jens Schmitt (TU Kaiserslautern, DE), Giovanni Stea (University of Pisa, IT)
March 3–8, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19101

19102 – 3D Morphable Models
Bernhard Egger (MIT – Cambridge, US), William Smith (University of York, GB), Christian Theobalt (MPI für Informatik – Saarbrücken, DE), Thomas Vetter (Universität Basel, CH)
March 3–8, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19102

19111 – Theoretical Foundations of Storage Systems
Martin Farach-Colton (Rutgers University – Piscataway, US), Inge Li Gørtz (Technical University of Denmark – Lyngby, DK), Rob Johnson (VMware – Palo Alto, US), Donald E. Porter (University of North Carolina at Chapel Hill, US)
March 10–15, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19111

19112 – Engineering Reliable Multiagent Systems
Jürgen Dix (TU Clausthal, DE), Brian Logan (University of Nottingham, GB), Michael Winikoff (University of Otago, NZ)
March 10–15, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19112

19121 – Computational Complexity of Discrete Problems
Anna Gál (University of Texas – Austin, US), Oded Regev (New York University, US), Rahul Santhanam (University of Oxford, GB), Till Tantau (Universität zu Lübeck, DE)
March 17–22, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19121

19131 – Algorithmic Problems in Group Theory
Volker Diekert (Universität Stuttgart, DE), Olga Kharlampovic (The City University of New York, US), Markus Lohrey (Universität Siegen, DE), Alexei Myasnikov (Stevens Institute of Technology – Hoboken, US)
March 24–29, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19131

19132 – Users and Automated Driving Systems: How Will We Interact with Tomorrow's Vehicles?
Susanne Boll (Universität Oldenburg, DE), Andrew Kun (University of New Hampshire – Durham, US), Andreas Riener (TH Ingolstadt, DE), C. Y. David Yang (AAA Foundation for Traffic Safety – Washington, US)
March 24–29, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19132

19141 – Programmable Network Data Planes
Gianni Antichi (Queen Mary University of London, GB), Theophilus Benson (Brown University – Providence, US), Nate Foster (Cornell University – Ithaca, US), Fernando M. V. Ramos (University of Lisbon, PT), Justine Sherry (Carnegie Mellon University – Pittsburgh, US)
March 31 to April 5, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19141

19151 – Visual Computing in Materials Sciences
Christoph Heinzl (FH Oberösterreich – Wels, AT), Robert Michael Kirby (University of Utah – Salt Lake City, US), Stepan V. Lomov (KU Leuven, BE), Guillermo Requena (DLR – Köln, DE), Rüdiger Westermann (TU München, DE)
April 7–12, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19151

19152 – Emerging Hardware Techniques and EDA Methodologies for Neuromorphic Computing
Krishnendu Chakrabarty (Duke University – Durham, US), Tsung-Yi Ho (National Tsing Hua University – Hsinchu, TW), Hai Li (Duke University – Durham, US), Ulf Schlichtmann (TU München, DE)
April 7–10, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19152

19171 – Ethics and Trust: Principles, Verification and Validation
Michael Fisher (University of Liverpool, GB), Christian List (London School of Economics, GB), Marija Slavkovik (University of Bergen, NO), Astrid Weiss (TU Wien, AT)
April 22–26, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19171

19172 – Computational Creativity Meets Digital Literary Studies
Tarek Richard Besold (Telefonica Research – Barcelona, ES), Pablo Gervás (Complutense University of Madrid, ES), Evelyn Gius (TU Darmstadt, DE), Sarah Schulz (Ada Health – Berlin, DE)
April 22–25, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19172

19181 – Computational Geometry
Siu-Wing Cheng (HKUST – Kowloon, HK), Anne Driemel (Universität Bonn, DE), Jeff Erickson (University of Illinois – Urbana-Champaign, US)
April 28 to May 3, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19181

19182 – Multi-Document Information Consolidation
Ido Dagan (Bar-Ilan University – Ramat Gan, IL), Iryna Gurevych (TU Darmstadt, DE), Dan Roth (University of Pennsylvania – Philadelphia, US), Amanda Stent (Bloomberg – New York, US)
April 28 to May 3, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19182

19191 – Software Evolution in Time and Space: Unifying Version and Variability Management
Thorsten Berger (Chalmers and University of Gothenburg, SE), Marsha Chechik (University of Toronto, CA), Timo Kehrer (HU Berlin, DE), Manuel Wimmer (Johannes Kepler Universität Linz, AT)
May 5–10, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19191

19192 – Visual Analytics for Sets over Time and Space
Sara Irina Fabrikant (Universität Zürich, CH), Silvia Miksch (TU Wien, AT), Alexander Wolff (Universität Würzburg, DE)
May 5–10, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19192

19202 – Approaches and Applications of Inductive Programming
Luc De Raedt (KU Leuven, BE), Richard Evans (Google DeepMind – London, GB), Stephen H. Muggleton (Imperial College London, GB), Ute Schmid (Universität Bamberg, DE)
May 12–17, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19202

19211 – Enumeration in Data Management
Endre Boros (Rutgers University – Piscataway, US), Benny Kimelfeld (Technion – Haifa, IL), Reinhard Pichler (TU Wien, AT), Nicole Schweikardt (HU Berlin, DE)
May 19–24, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19211

19212 – Topology, Computation and Data Analysis
Michael Kerber (TU Graz, AT), Vijay Natarajan (Indian Institute of Science – Bangalore, IN), Bei Wang (University of Utah – Salt Lake City, US)
May 19–24, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19212

19222 – Control of Networked Cyber-Physical Systems
John S. Baras (University of Maryland – College Park, US), Sandra Hirche (TU München, DE), Kay Römer (TU Graz, AT), Klaus Wehrle (RWTH Aachen, DE)
May 26–29, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19222

19231 – Empirical Evaluation of Secure Development Processes
Adam Shostack (Seattle, US), Matthew Smith (Universität Bonn and Fraunhofer FKIE, DE), Sam Weber (Carnegie Mellon University – Pittsburgh, US), Mary Ellen Zurko (MIT Lincoln Laboratory – Lexington, US)
June 2–7, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19231

19232 – Ubiquitous Computing Education: Why, What, and How
Audrey Girouard (Carleton University – Ottawa, CA), Andrew Kun (University of New Hampshire – Durham, US), Anne Roudaut (University of Bristol, GB), Orit Shaer (Wellesley College, US)
June 2–7, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19232

19241 – 25 Years of the Burrows-Wheeler Transform
Travis Gagie (Universidad Diego Portales, CL), Giovanni Manzini (University of Eastern Piedmont – Alessandria, IT), Gonzalo Navarro (University of Chile – Santiago de Chile, CL), Jens Stoye (Universität Bielefeld, DE)
June 10–14, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19241

19261 – Distributed Computing with Permissioned Blockchains and Databases
C. Mohan (IBM Almaden Center – San Jose, US), Beng Chin Ooi (National University of Singapore, SG), Andreas Reuter (Heidelberg Laureate Forum Foundation, DE), Gottfried Vossen (Universität Münster, DE)
June 23–28, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19261

19262 – Astrographics: Interactive Data-Driven Journeys through Space
Alyssa A. Goodman (Harvard-Smithsonian Center for Astrophysics, US), Charles D. Hansen (University of Utah – Salt Lake City, US), Daniel Weiskopf (Universität Stuttgart, DE), Anders Ynnerman (Linköping University, SE)
June 23–26, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19262

19271 – Graph Colouring: From Structure to Algorithms
Maria Chudnovsky (Princeton University, US), Daniel Paulusma (Durham University, GB), Oliver Schaudt (RWTH Aachen, DE)
June 30 to July 5, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19271

19272 – Real VR – Importing the Real World into Immersive VR and Optimizing the Perceptual Experience of Head-Mounted Displays
Marcus A. Magnor (TU Braunschweig, DE), Alexander Sorkine-Hornung (Oculus VR – Zürich, CH)
June 30 to July 3, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19272

19281 – Notional Machines and Programming Language Semantics in Education
Mark Guzdial (University of Michigan – Ann Arbor, US), Shriram Krishnamurthi (Brown University – Providence, US), Juha Sorva (Aalto University, FI), Jan Vahrenhold (Universität Münster, DE)
July 7–12, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19281

19282 – Data Series Management
Anthony Bagnall (University of East Anglia – Norwich, GB), Richard L. Cole (Tableau Software – Palo Alto, US), Themis Palpanas (Paris Descartes University, FR), Konstantinos Zoumpatianos (Harvard University – Cambridge, US)
July 7–12, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19282

19291 – Values in Computing
Christoph Becker (University of Toronto, CA), Gregor Engels (Universität Paderborn, DE), Andrew Feenberg (Simon Fraser University – Burnaby, CA), Maria Angela Ferrario (Lancaster University, GB), Geraldine Fitzpatrick (TU Wien, AT)
July 14–19, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19291

19292 – Mobile Data Visualization
Eun Kyoung Choe (University of Maryland – College Park, US), Raimund Dachselt (TU Dresden, DE), Petra Isenberg (INRIA Saclay – Orsay, FR), Bongshin Lee (Microsoft Research – Redmond, US)
July 14–19, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19292

19301 – Secure Composition for Hardware Systems
Divya Arora (Intel – Santa Clara, US), Ilia Polian (Universität Stuttgart, DE), Francesco Regazzoni (University of Lugano, CH), Patrick Schaumont (Virginia Polytechnic Institute – Blacksburg, US)
July 21–26, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19301

19302 – Cybersafety Threats – from Deception to Aggression
Zinaida Benenson (Universität Erlangen-Nürnberg, DE), Marianne Junger (University of Twente, NL), Daniela Oliveira (University of Florida – Gainesville, US), Gianluca Stringhini (Boston University, US)
July 21–26, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19302

19331 – Software Protection Decision Support and Evaluation Methodologies
Christian Collberg (University of Arizona – Tucson, US), Mila Dalla Preda (University of Verona, IT), Bjorn De Sutter (Ghent University, BE), Brecht Wyseur (Kudelski Group – Cheseaux, CH)
August 11–16, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19331

19341 – Algorithms and Complexity for Continuous Problems
Dmitriy Bilyk (University of Minnesota – Minneapolis, US), Aicke Hinrichs (Johannes Kepler Universität Linz, AT), Frances Y. Kuo (UNSW Sydney, AU), Klaus Ritter (TU Kaiserslautern, DE)
August 18–23, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19341

19342 – Advances and Challenges in 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)
August 18–23, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19342

19351 – Computational Proteomics
Nuno Bandeira (University of California – San Diego, US), Ileana M. Cristea (Princeton University, US), Lennart Martens (Ghent University, BE)
August 25–30, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19351

19352 – Computation in Low-Dimensional Geometry and Topology
Maarten Löffler (Utrecht University, NL), Anna Lubiw (University of Waterloo, CA), Saul Schleimer (University of Warwick – Coventry, GB), Erin Moriarty Wolf Chambers (St. Louis University, US)
August 25–30, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19352

19361 – Logic and Learning
Michael Benedikt (University of Oxford, GB), Kristian Kersting (TU Darmstadt, DE), Phokion G. Kolaitis (University of California – Santa Cruz & IBM Almaden Research Center – San Jose, US), Daniel Neider (MPI-SWS – Kaiserslautern, DE)
September 1–6, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19361

19371 – Deduction Beyond Satisfiability
Carsten Fuhs (Birkbeck, University of London, GB), Philipp Rümmer (Uppsala University, SE), Renate Schmidt (University of Manchester, GB), Cesare Tinelli (University of Iowa – Iowa City, US)
September 8–13, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19371

19381 – Application-Oriented Computational Social Choice
Umberto Grandi (University Toulouse Capitole, FR), Stefan Napel (Universität Bayreuth, DE), Rolf Niedermeier (TU Berlin, DE), Kristen Brent Venable (IHMC – Pensacola, US)
September 15–20, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19381

19391 – Data Ecosystems: Sovereign Data Exchange among Organizations
Cinzia Cappiello (Polytechnic University of Milan, IT), Avigdor Gal (Technion – Haifa, IL), Matthias Jarke (RWTH Aachen, DE), Jakob Rehof (TU Dortmund, DE)
September 22–27, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19391

19401 – Comparative Theory for Graph Polynomials
Jo Ellis-Monaghan (Saint Michael's College – Colchester, US), Andrew Goodall (Charles University – Prague, CZ), Iain Moffatt (Royal Holloway University of London, GB), Kerri Morgan (Deakin University – Melbourne, AU)
September 29 to October 4, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19401

19411 – Social Agents for Teamwork and Group Interactions
Elisabeth André (Universität Augsburg, DE), Ana Paiva (INESC-ID – Porto Salvo, PT), Julie Shah (MIT – Cambridge, US), Selma Šabanovic (Indiana University – Bloomington, US)
October 6–11, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19411

19421 – Quantum Cryptanalysis
Michele Mosca (University of Waterloo, CA), Maria Naya-Plasencia (INRIA – Paris, FR), Rainer Steinwandt (Florida Atlantic University – Boca Raton, US), Krysta Svore (Microsoft Corporation – Redmond, US)
October 13–18, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19421

19431 – Theory of Randomized Optimization Heuristics
Carola Doerr (Sorbonne University – Paris, FR), Carlos M. Fonseca (University of Coimbra, PT), Tobias Friedrich (Hasso-Plattner-Institut – Potsdam, DE), Xin Yao (Southern Univ. of Science and Technology – Shenzen, CN)
October 20–25, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19431

19432 – Analysis of Autonomous Mobile Collectives in Complex Physical Environments
Mario Gleirscher (University of York, GB), Anne E. Haxthausen (Technical University of Denmark – Lyngby, DK), Martin Leucker (Universität Lübeck, DE), Sven Linker (University of Liverpool, GB)
October 20–23, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19432

19442 – Programming Languages for Distributed Systems and Distributed Data Management
Carla Ferreira (New University of Lisbon, PT), Philipp Haller (KTH Royal Institute of Technology – Stockholm, SE), Volker Markl (TU Berlin, DE), Guido Salvaneschi (TU Darmstadt, DE), Cristina Videira Lopes (University of California – Irvine, US)
October 27–31, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19442

19443 – Algorithms and Complexity in Phylogenetics
Magnus Bordewich (Durham University, GB), Britta Dorn (Universität Tübingen, DE), Simone Linz (University of Auckland, NZ), Rolf Niedermeier (TU Berlin, DE)
October 27–31, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19443

19451 – Biggest Failures in Security
Frederik Armknecht (Universität Mannheim, DE), Ingrid Verbauwhede (KU Leuven, BE), Melanie Volkamer (KIT – Karlsruher Institut für Technologie, DE), Moti Yung (Columbia University – New York, US)
November 3–8, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19451

19452 – Machine Learning Meets Visualization to Make Artificial Intelligence Interpretable
Enrico Bertini (NYU – Brooklyn, US), Peer-Timo Bremer (LLNL – Livermore, US), Daniela Oelke (Siemens AG – München, DE), Jayaraman Thiagarajan (LLNL – Livermore, US)
November 3–8, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19452

19461 – Conversational Search
Avishek Anand (Leibniz Universität Hannover, DE), Lawrence Cavedon (RMIT University – Melbourne, AU), Hideo Joho (University of Tsukuba – Ibaraki, JP), Mark Sanderson (RMIT University – Melbourne, AU), Benno Stein (Bauhaus-Universität Weimar, DE)
November 10–15, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19461

19471 – BOTse: Bots in Software Engineering
James D. Herbsleb (Carnegie Mellon University – Pittsburgh, US), Carolyn Penstein Rosé (Carnegie Mellon University – Pittsburgh, US), Alexander Serebrenik (TU Eindhoven, NL), Margaret-Anne Storey (University of Victoria, CA), Thomas Zimmermann (Microsoft Corporation – Redmond, US)
November 17–22, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19471

19481 – Composing Model-Based Analysis Tools
Francisco Durán (University of Málaga, ES), Robert Heinrich (KIT – Karlsruhe, DE), Diego Pérez-Palacín (Linnaeus University – Växjö, SE), Carolyn L. Talcott (SRI – Menlo Park, US), Steffen Zschaler (King's College London, GB)
November 24–29, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19481

19491 – Big Graph Processing Systems
Angela Bonifati (University Claude Bernard – Lyon, FR), Alexandru Iosup (VU University Amsterdam, NL), Sherif Sakr (University of Tartu, EE), Hannes Voigt (Neo4j – Leipzig, DE)
December 1–6, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19491

19502 – Future Automotive HW/SW Platform Design
Xiaobo Sharon Hu (University of Notre Dame, US), Selma Saidi (TU Dortmund, DE), Sebastian Steinhorst (TU München, DE), Dirk Ziegenbein (Robert Bosch GmbH – Stuttgart, DE)
December 8–11, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19502

19511 – Artificial and Computational Intelligence in Games: Revolutions in Computational Game AI
Jialin Liu (Southern Univ. of Science and Technology – Shenzen, CN), Tom Schaul (Google DeepMind – London, GB), Pieter Spronck (Tilburg University, NL), Julian Togelius (New York University, US)
December 15–20, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19511

19512 – Interactive Design and Simulation
Thomas A. Grandine (The Boeing Company – Seattle, US), Jörg Peters (University of Florida – Gainesville, US), Ulrich Reif (TU Darmstadt, DE), Olga Sorkine-Hornung (ETH Zürich, CH)
December 15–20, 2019 | Dagstuhl Seminar | https://www.dagstuhl.de/19512

Dagstuhl-Perspektiven-Workshops 14.2 *Dagstuhl Perspectives Workshops*

19072 – The Role of Non-Monotonic Reasoning in Future Development of Artificial Intelligence
Anthony Hunter (University College London, GB), Gabriele Kern-Isberner (TU Dortmund, DE), Thomas Meyer (University of Cape Town, ZA), Renata Wassermann (University of Sao Paulo, BR)
February 10–15, 2019 | Dagstuhl Perspectives Workshop | https://www.dagstuhl.de/19072

19482 – Diversity, Fairness, and Data-Driven Personalization in (News) Recommender System
Abraham Bernstein (Universität Zürich, CH), Claes De Vreese (University of Amsterdam, NL), Natali Helberger (University of Amsterdam, NL), Wolfgang Schulz (Universität Hamburg, DE), Katharina A. Zweig (TU Kaiserslautern, DE)
November 24–29, 2019 | Dagstuhl Perspectives Workshop | https://www.dagstuhl.de/19482

GI-Dagstuhl-Seminare 14.3 *GI-Dagstuhl Seminars*

19023 – Explainable Software for Cyber-Physical Systems
Joel Greenyer (Leibniz Universität Hannover, DE), Malte Lochau (TU Darmstadt, DE), Thomas Vogel (HU Berlin, DE)
January 6–11, 2019 | GI-Dagstuhl Seminar | https://www.dagstuhl.de/19023

19363 – Algorithms for Big Data
Timo Bingmann (KIT – Karlsruher Institut für Technologie, DE), Tim Conrad (FU Berlin, DE), Ulrich Carsten Meyer (Goethe-Universität – Frankfurt am Main, DE), Matthias Mnich (Universität Bonn, DE)
September 1–6, 2019 | GI-Dagstuhl Seminar | https://www.dagstuhl.de/19363

Lehrveranstaltungen 14.4 *Educational Events*

19093 – 2nd Winter School on Operating Systems (WSOS 2019) Focus: "Operating Systems in Research and Industry"
Marcel Carsten Baunach (TU Graz, AT), Michael Engel (Hochschule Coburg, DE), Dieter Kasper (Fujitsu – München, DE), Olaf Spinczyk (Universität Osnabrück, DE)
February 24 to March 1, 2019 | Educational Event | https://www.dagstuhl.de/19093

19123 – Summer School "Data Management Techniques"
Goetz Graefe (Google – Madison, US)
March 17–21, 2019 | Educational Event | https://www.dagstuhl.de/19123

19203 – 3rd Summer Datathon on Linguistic Linked Open Data (SD-LLOD 2019)
Christian Chiarcos (Goethe-Universität – Frankfurt am Main, DE), Jorge Gracia (University of Zaragoza, ES), John McCrae (National University of Ireland – Galway, IE)
May 12–17, 2019 | Educational Event | https://www.dagstuhl.de/19203

19333 – Summer School "Metaprogramming"
Yukiyoshi Kameyama (University of Tsukuba, JP), Ohad Kammar (University of Edinburgh, GB), Jeremy Yallop (University of Cambridge, GB)
August 11–16, 2019 | Educational Event | https://www.dagstuhl.de/19333

19383 – Summer School "Development, Deployment, and Runtime of Context-Aware Software Systems"
Wolfgang Lehner (TU Dresden, DE)
September 15–18, 2019 | Educational Event | https://www.dagstuhl.de/19383

19393 – Autumn School 2019 for Information Retrieval and Information Foraging
Ingo Frommholz (University of Bedfordshire – Luton, GB), Norbert Fuhr (Universität Duisburg-Essen, DE), Ralf Schenkel (Universität Trier, DE)
September 22–27, 2019 | Educational Event | https://www.dagstuhl.de/19393

19503 – Lehrerfortbildung in Informatik
Michael Gerke (Schloss Dagstuhl – Saarbrücken, DE), Gerrit Müller (Peter-Wust-Gymnasium – Merzig, DE & LPM Saarbrücken, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 11–13, 2019 | Educational Event | https://www.dagstuhl.de/19503

Forschungsgruppentreffen 14.5 *Research Group Meetings*

19039 – Forschungsgast
Guillaume Aucher (University of Rennes 1 & IRISA Rennes, FR)
January 17–19, 2019 | Research Group Meeting | https://www.dagstuhl.de/19039

19143 – Software Engineering Forschungsmethoden Training
Sven Apel (Universität Passau, DE), Eric Bodden (Universität Paderborn, DE), Lars Grunske (HU Berlin, DE)
March 31 to April 3, 2019 | Research Group Meeting | https://www.dagstuhl.de/19143

19144 – Klausurtagung Telematik Karlsruhe
Robert Bauer (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)
April 3–5, 2019 | Research Group Meeting | https://www.dagstuhl.de/19144

19145 – Lehrstuhltreffen Rechtsinformatik
Christoph Sorge (Universität des Saarlandes, DE)
April 3–5, 2019 | Research Group Meeting | https://www.dagstuhl.de/19145

19153 – Digital Archaeoludology
Cameron Browne (Maastricht University, NL)
April 10–12, 2019 | Research Group Meeting | https://www.dagstuhl.de/19153

19163 – GIBU 2019: GI-Beirat der Universitätsprofessoren
Lars Grunske (HU Berlin, DE)
April 14–16, 2019 | Research Group Meeting | https://www.dagstuhl.de/19163

19164 – Lehrstuhltreffen "Embedded Intelligence"
Bernhard Sick (Universität Kassel, DE)
April 14–17, 2019 | Research Group Meeting | https://www.dagstuhl.de/19164

19173 – Artificial Intelligence for Production Technology – Data Products, Smart Services and Resilience
Wolfgang Maaß (Universität des Saarlandes – Saarbrücken, DE)
April 25–26, 2019 | Research Group Meeting | https://www.dagstuhl.de/19173

19204 – Aufbereitung der Geschichte der GI
Stefan Jähnichen (TU Berlin, DE)
May 16–18, 2019 | Research Group Meeting | https://www.dagstuhl.de/19204

19223 – Kolloquium zum GI Dissertationspreis 2018
Steffen Hölldobler (TU Dresden, DE)
May 26–29, 2019 | Research Group Meeting | https://www.dagstuhl.de/19223

19243 – Workshop Buchprojekt "Applied Machine Intelligence"
Thomas Hoppe (Fraunhofer FOKUS – Berlin, DE), Bernhard Humm (Hochschule Darmstadt, DE), Anatol Reibold (DDG: IX GmbH – Darmstadt, DE)
June 10–14, 2019 | Research Group Meeting | https://www.dagstuhl.de/19243

19253 – Gemeinsamer Workshop der Graduiertenkollegs GRK 2050 und GRK 2236
Helen Bolke-Hermanns (RWTH Aachen, DE), Joost-Pieter Katoen (RWTH Aachen, DE), Max Mühlhäuser (TU Darmstadt, DE)
June 16–19, 2019 | Research Group Meeting | https://www.dagstuhl.de/19253

19269 – Forschungsaufenthalt
Christoph Becker (University of Toronto, CA)
June 23 to July 14, 2019 | Research Group Meeting | https://www.dagstuhl.de/19269

19299 – Forschungsaufenthalt
Iris van Rooij (Radboud University Nijmegen, NL)
July 14–20, 2019 | Research Group Meeting | https://www.dagstuhl.de/19299

19309 – Forschungsaufenthalt
Christoph Becker (University of Toronto, CA)
July 20–25, 2019 | Research Group Meeting | https://www.dagstuhl.de/19309

19373 – Dagstuhl Workshop "Intelligent Methods for Test and Reliability"
Ilia Polian (Universität Stuttgart, DE)
September 11–13, 2019 | Research Group Meeting | https://www.dagstuhl.de/19373

19384 – Klausurtagung AG Robotersysteme TU KL
Karsten Berns (TU Kaiserslautern, DE)
September 19–20, 2019 | Research Group Meeting | https://www.dagstuhl.de/19384

19403 – DDI 4 Core – Development of a Robust and Sustainable Model
Arofan Gregory (Jaffrey, US), Hilde Orten (NSD – Bergen, NO), Joachim Wackerow (GESIS – Mannheim, DE)
September 29 to October 4, 2019 | Research Group Meeting | https://www.dagstuhl.de/19403

19413 – Interoperability of Metadata Standards in Cross-Domain Science, Health, and Social Science Applications II
Simon Cox (CSIRO – Clayton South, AU), Arofan Gregory (Jaffrey, US), Simon Hodson (CODATA – Paris, FR), Steven McEachern (Australian National University – Acton, AU), Joachim Wackerow (GESIS – Mannheim, DE)
October 6–11, 2019 | Research Group Meeting | https://www.dagstuhl.de/19413

19419 – Forschungsaufenthalt
Danny Weyns (KU Leuven, BE)
October 6–12, 2019 | Research Group Meeting | https://www.dagstuhl.de/19419

19423 – OpenML Hackathon
Heidi Seibold (LMU München, DE), Joaquin Vanschoren (TU Eindhoven, NL)
October 13–18, 2019 | Research Group Meeting | https://www.dagstuhl.de/19423

19433 – Klausurtagung Wolpert/Schömer
Nicola Wolpert (University of Applied Sciences – Stuttgart, DE)
October 23–25, 2019 | Research Group Meeting | https://www.dagstuhl.de/19433

19434 – Lehrstuhltreffen AG Zeller
Andreas Zeller (CISPA – Saarbrücken, DE)
October 23–25, 2019 | Research Group Meeting | https://www.dagstuhl.de/19434

19464 – Arbeitstreffen Text-Technology Lab
Giuseppe Abrami (Goethe-Universität – Frankfurt am Main, DE), Alexander Mehler (Goethe-Universität – Frankfurt am Main, DE)
November 11–13, 2019 | Research Group Meeting | https://www.dagstuhl.de/19464

19473 – Retreat of the Research Training Group "Adaptive Information Preparation from Heterogeneous Sources" (AIPHES)
Federico López (HITS – Heidelberg, DE), Michael Strube (HITS – Heidelberg, DE)
November 18–20, 2019 | Research Group Meeting | https://www.dagstuhl.de/19473

19474 – Secan Lab Seminar
Thomas Engel (University of Luxembourg, LU)
November 21–22, 2019 | Research Group Meeting | https://www.dagstuhl.de/19474

19499 – Forschungsaufenthalt
Anette Frank (Universität Heidelberg, DE)
December 4–6, 2019 | Research Group Meeting | https://www.dagstuhl.de/19499

19504 – GI Tagung Roboter Steuerungsarchitekturen
Karsten Berns (TU Kaiserslautern, DE)
December 8–10, 2019 | Research Group Meeting | https://www.dagstuhl.de/19504

19505 – inf-schule.de: Weiterentwicklung des elektronischen und interaktiven Lehrbuches für den Informatikunterricht
Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 10–11, 2019 | Research Group Meeting | https://www.dagstuhl.de/19505

Leibniz
Gemeinschaft

Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2019 | ISSN 2199-1995
http://www.dagstuhl.de