SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2021

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2021

Herausgeber	Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Germany Telefon: +49 681 302 – 4396 Fax: +49 681 302 – 4397 E-Mail: service@dagstuhl.de
Registernummer	Amtsgericht Saarbrücken HRB 63800
Vorsitzender des Aufsichtsrates	Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen
Geschäftsführung	Prof. Raimund Seidel, Ph. D. Heike Meißner
Gesellschafter	Gesellschaft für Informatik e. V., Deutschland Universität des Saarlandes, Deutschland Technische Universität Kaiserslautern, Deutschland Karlsruher Institut für Technologie (KIT), Deutschland Technische Universität Darmstadt, Deutschland Universität Stuttgart, Deutschland Universität Trier, Deutschland Johann Wolfgang Goethe-Universität Frankfurt am Main, Deutschland Institut National de Recherche en Informatique et en Automatique (INRIA), Frankreich Centrum voor Wiskunde en Informatica (CWI), Niederlande Max-Planck-Gesellschaft zur Förderung der Wissenschaften e. V., Deutschland
Einbindung	Mitglied der Leibniz-Gemeinschaft Mitglied im Leibniz-Forschungsverbund Science 2.0
Verantwortlicher Redakteur	Prof. Raimund Seidel, Ph. D.
Redaktion	Dr. Michael Wagner
Mitarbeit	Dr. Marcel R. Ackermann, Heike Clemens, Dr. Michael Didas, Dr. Andreas Dolzmann, Jutka Gasiorowski, Dr.-Ing. Michael Gerke, Dagmar Glaser, Oliver Hoffmann, Dr. Marsha Kleinbauer, Heike Meißner, Petra Meyer, Dr. Florian Reitz, Simone Schilke, Thomas Schillo, Christina Schwarz, Prof. 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	September 2022
Copyright	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Lizenz	Creative Commons BY-ND 4.0

Leibniz Gemeinschaft

Zertifikat seit 2012 audit berufundfamilie

Vorwort
Foreword

Auch das Jahr 2021 war für Schloss Dagstuhl noch stark von der Covid-19 Pandemie geprägt. Reisebeschränkungen aber auch Pandemievorschriften ließen manchmal gar keinen Betrieb vor Ort zu. Seminare wurden teilweise abgesagt und zur raschen Wiederbeantragung ermuntert, einige fanden „virtuell“ nur mit Fernteilnehmern statt und wieder andere in einem hybriden Modus mit nur wenigen Teilnehmern vor Ort und den anderen zugeschaltet. Erst in der zweiten Jahreshälfte näherten wir uns einem „Normalbetrieb“ an mit den meisten Teilnehmern vor Ort, wenn auch bei einer leicht beschränkten Teilnehmerzahl.

Diese Zeit war und ist immer noch ein große Herausforderung für unsere Mitarbeiter: die vielen kurzfristigen Entscheidungen und Anpassungen und die sich stets ändernden Vorschriften und Regeln, die typischerweise ausländische Gäste nicht so sehr im Augenmerk hatten.

Wir konnten die Zeit allerdings auch positiv nutzen. So wurde zum Beispiel der große Vortragssaal renoviert und umgestaltet und seine Akustik deutlich verbessert.

Eines hat uns die Pandemie aber deutlich gezeigt und wird auch von unseren Gästen immer wieder betont: wie wichtig und produktiv die persönliche Begegnung in unserem Wissenschaftszentrum ist.

For Schloss Dagstuhl, the year 2021 was still strongly influenced by the Covid-19 pandemic. At times, travel restrictions but also pandemic regulations did not allow any meetings on site. Some seminars were cancelled and speedy re-application encouraged, some were held “virtually” with only remote participants, and still others in a hybrid mode with only a few participants on site and the others connected via zoom. It was not until the second half of the year that we approached a “normal mode” with most participants on site, albeit with a slightly limited number of participants.

This all was, and still is, a great challenge for our staff: the many last-minute decisions and adjustments and the ever-changing rules and regulations, which typically did not have foreign guests so much in mind.

We were able to make positive use of the time. For example, the large lecture hall has been renovated and redesigned and its acoustics significantly improved.

The pandemic has, however, clearly shown us one thing, which is also emphasized again and again by our guests: the importance and usefulness of personal encounters and interactions in our science center.

Im Namen der Geschäftsleitung

Prof. Raimund Seidel, Ph. D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

On behalf of the Managing Directors

Inhaltsverzeichnis
Contents

Vorwort
Foreword . . . iii

1 Das Zentrum Schloss Dagstuhl
Schloss Dagstuhl Center **1**
1.1 Dagstuhls Leitbild
Dagstuhl's Mission . . . 2
1.2 Neuigkeiten in 2021
News from 2021 . . . 4

2 Seminare und Workshops
Seminars and Workshops **9**
2.1 Dagstuhl-Seminare
Dagstuhl Seminars . . . 10
2.2 Dagstuhl-Perspektiven-Workshops
Dagstuhl Perspectives Workshops . . . 11
2.3 Einreichung der Anträge und Begutachtungsverfahren
Proposal Submission and Review Process . . . 11
2.4 Seminar-Programm 2021
Seminar-Programm in 2021 . . . 12
2.5 Angaben zu Teilnehmern und Organisatoren
Participant and Organizer Data . . . 13
2.6 Themen und Forschungsgebiete
Topics and Research Areas . . . 14
2.7 Weitere Veranstaltungstypen
Further Event Types . . . 17
2.8 Qualitätssicherung
Quality Assurance . . . 17
2.9 Auslastung des Zentrums
Utilization of the Center . . . 18

3 Bibliographiedatenbank dblp
dblp computer science bibliography **21**
3.1 Offene Bibliographiedaten für die Informatik
Open Bibliographic Data in Computer Science . . . 22
3.2 Schloss Dagstuhl und dblp
Schloss Dagstuhl and dblp . . . 23
3.3 Statistiken der Datenakquise
Data Acquisition Statistics . . . 24
3.4 Statistiken der Datenkuration
Data Curation Statistics . . . 24
3.5 Nutzungsstatistiken
Usage Statistics . . . 26

4 Dagstuhl Publishing
Dagstuhl Publishing **27**
4.1 Portfolio
Portfolio . . . 28
4.2 Infrastruktur
Infrastructure . . . 35

5 Resonanz
Feedback **37**
5.1 Resonanz zu Seminaren und Workshops
Feedback on Seminars and Workshops . . . 38
5.2 Resonanz zur Bibliographiedatenbank dblp
Feedback on the dblp computer science bibliography . . . 42
5.3 Resonanz zu Dagstuhl Publishing
Feedback on Dagstuhl Publishing . . . 44

6 Die Seminare in 2021
The 2021 Seminars **45**
6.1 Scalable Data Structures . . . 48
6.2 Autonomous Agents on the Web . . . 50
6.3 Computational Complexity of Discrete Problems . . . 54
6.4 Multi-Level Graph Representation for Big Data Arising in Science Mapping . . . 56
6.5 Temporal Graphs: Structure, Algorithms, Applications . . . 58
6.6 Computational Geometry . . . 60
6.7 Approaches and Applications of Inductive Programming . . . 62
6.8 Serverless Computing . . . 63
6.9 Transparency by Design . . . 68
6.10 Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments . . . 70
6.11 Compute-First Networking . . . 77
6.12 Quantum Complexity: Theory and Application . . . 78
6.13 Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users . . . 80
6.14 Computational Proteomics . . . 82
6.15 Towards Climate-Friendly Internet Research . . . 84
6.16 Data Structures for Modern Memory and Storage Hierarchies . . . 85
6.17 Scalable Handling of Effects . . . 86
6.18 Parameterized Complexity in Graph Drawing . . . 88
6.19 Matching Under Preferences: Theory and Practice . . . 92
6.20 Approximate Systems . . . 94
6.21 Coalition Formation Games . . . 96
6.22 Understanding I/O Behavior in Scientific and Data-Intensive Computing . . . 98
6.23 Identifying Key Enablers in Edge Intelligence . . . 100
6.24 Universals of Linguistic Idiosyncrasy in Multilingual Computational Linguistics . . . 102
6.25 Higher-Order Graph Models: From Theoretical Foundations to Machine Learning . . . 104
6.26 Extending the Synergies Between SAT and Description Logics . . . 108
6.27 Structure and Learning . . . 110
6.28 Integrated Deduction . . . 112
6.29 Behavioural Types: Bridging Theory and Practice . . . 114
6.30 Conversational Agent as Trustworthy Autonomous System (Trust-CA) . . . 116
6.31 Sparsity in Algorithms, Combinatorics and Logic . . . 118
6.32 Visualization of Biological Data – From Analysis to Communication . . . 120
6.33 Digital Disinformation: Taxonomy, Impact, Mitigation, and Regulation . . . 122
6.34 Machine Learning in Sports . . . 124
6.35 Quantum Cryptanalysis . . . 126
6.36 Rigorous Methods for Smart Contracts . . . 128
6.37 Probabilistic Numerical Methods – From Theory to Implementation . . . 130
6.38 Adaptive Resource Management for HPC Systems . . . 132
6.39 Ensuring the Reliability and Robustness of Database Management Systems . . . 134
6.40 Managing Industrial Control Systems Security Risks for Cyber Insurance . . . 136
6.41 Unambiguity in Automata Theory . . . 140
6.42 Descriptive Set Theory and Computable Topology . . . 142

6.43 Foundations of Persistent Programming . . . 146
6.44 Geometric Modeling: Interoperability and New Challenges . . . 148
6.45 Geometric Logic, Constructivisation, and Automated Theorem Proving . . . 149
6.46 Secure Compilation . . . 150
6.47 Representing and Solving Spatial Problems . . . 152

7 Öffentlichkeitsarbeit
Public Relations and Outreach **157**
7.1 Pressemitteilungen und Medienarbeit
Press Releases and Media Work . . . 158

8 Einrichtungen
Facilities **159**
8.1 Hauptstandort in Wadern
Main Site in Wadern . . . 160
8.2 Geschäftsstelle in Saarbrücken
Dagstuhl Office in Saarbrücken . . . 163
8.3 Dagstuhl an der Universität Trier
Dagstuhl at the University of Trier . . . 163

9 Zentrale Dienste
Central Services **165**
9.1 Bibliothek
Research Library . . . 166
9.2 IT-Service
IT Service . . . 168

10 Kunst
Art **169**
10.1 Dagstuhl als Galerie
Dagstuhl as an Art Gallery . . . 170
10.2 Dagstuhls permanente Kunstausstellung
Dagstuhl's Permanent Art Exhibition . . . 170

11 Struktur der Gesellschaft
Structure of the Company **171**
11.1 Gründung und Gesellschafter
Formation and Shareholders . . . 172
11.2 Organe der Gesellschaft
Organs of the Organization . . . 172
11.3 Gremien der Gesellschaft
Bodies of the Organization . . . 174

12 Holger Hermanns, Erich Reindel: Förderverein „Freunde von Dagstuhl"
Holger Hermanns, Erich Reindel: Association "Friends of Dagstuhl" **179**

13 Statistiken
Statistics **183**
13.1 Statistiken zu Seminaren und Workshops
Statistics on Seminars and Workshops . . . 184
13.2 Statistiken zur Bibliographiedatenbank dblp
Statistics of the dblp computer science bibliography . . . 196
13.3 Statistiken zu Dagstuhl Publishing
Statistics of Dagstuhl Publishing . . . 200

14 Veranstaltungen 2021
Schedule of Events 2021 **203**
14.1 Dagstuhl-Seminare
Dagstuhl Seminars . . . 204

14.2 Dagstuhl-Perspektiven-Workshops
Dagstuhl Perspectives Workshops 207
14.3 GI-Dagstuhl-Seminare
GI-Dagstuhl Seminars 207
14.4 Lehrveranstaltungen
Educational Events 207
14.5 Forschungsgruppentreffen
Research Group Meetings 207

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) normalerweise 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 was growing 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) normally 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 has been 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 internatio-

Seminar and Workshop Program

The Dagstuhl Seminars and Dagstuhl Perspectives Workshops form the focus of the center's work. Whereas about 30 or 45 established and young researchers gather at

[1] Stand Januar 2021.
As of January 2021.

Fig. 1.1
Aerial photograph of Schloss Dagstuhl.

nale Forscher treffen sich eine halbe bis ganze Woche auf Schloss Dagstuhl, um im Rahmen eines Dagstuhl-Seminars 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.

the Dagstuhl Seminars to report on and discuss their current work, smaller groups of about 30 of the international elite of a field gather at the Dagstuhl Perspectives Workshops for the purpose of reflecting on the current status of research and potential development perspectives.

The Dagstuhl Seminars and Perspectives Workshops 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 built 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 5,4 Millionen Publikationseinträgen ist dblp die weltweit größte offene Sammlung bibliographischer Daten in der Informatik. Der dblp-Dienst ist darauf ausgerichtet, Forscher bei ihrer täglichen Arbeit zu unterstützen, etwa bei der Literaturrecherche oder beim Bezug von elektronisch verfügbaren Volltexten. Dabei gilt dblp in der Informatik insbesondere als die Referenzdatenbank für qualitätsgesicherte, normierte Bibliographiedaten. Aber auch Forschungsförderer und Entscheidungsträger unterstützt dblp, etwa durch das Pflegen und öffentlich Verfügbarmachen von personalisierten Publikationsnachweisen. Durch den Betrieb von dblp leistet Schloss Dagstuhl einen weiteren Beitrag im Rahmen seiner Mission zur Förderung der Erkennung, Verbreitung und Umsetzung neuer Informatikerkenntnisse auf international anerkanntem Niveau.

Details über dblp finden sich in Kapitel 3.

Dagstuhl Publishing

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

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

Kapitel 4 stellt Dagstuhls Verlagswesen ausführlicher dar.

dblp computer science bibliography

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

Listing more than 5.4 million articles, dblp is the world's most comprehensive open data collection of computer science research articles. The goal of dblp is to support computer scientists in their daily work, for example when reviewing the literature of a given author or subject area, or when searching for online full-text versions of research articles. The dblp database is often considered to be the reference database for quality-assured and normalized bibliographic metadata in computer science. Additionally, dblp supports funding agencies and decision makers by providing and curating personalized author 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.

Neuigkeiten in 2021 — 1.2 — *News from 2021*

Pandemie

Während die Bereiche Publishing und dblp kaum von der anhaltenden Covid-19 Pandemie betroffen waren und, ohne weitere Einschränkungen, vorwiegend im Homeoffice arbeiteten, war das Seminarwesen durch die vielen Auflagen und Reisebeschränkungen insbesondere für internationale Gäste weiterhin stark betroffen. Näheres zum Seminarwesen und den Einschränkungen dort im Abschnitt „Seminare und Workshops“ weiter unten. Auch

Pandemic

While the publishing and dblp departments were hardly affected by the Covid-19 pandemic and worked without any further restrictions – mainly in their home offices – the seminar department was again strongly affected by the many requirements and travel restrictions, especially for international guests. More details on the seminar department and the restrictions affecting it are given in section “Seminars and Workshops” below. The teacher

die Lehrerfortbildung, die Reihe „Dagstuhler Gespräche“ sowie teilweise die Kooperation mit dem Heidelberg Laureate Forum mussten auch für 2021 ausgesetzt werden, sollen aber wieder vollumfänglich aufgenommen werden, sobald es die Lage zulässt.

Das Team

Am Ende des Jahres 2021 beschäftigte Schloss Dagstuhl insgesamt 38,28 Vollzeitäquivalente bzw. 49 Mitarbeiter. Das LZI hat in 2021 eine Auszubildende und eine Mitarbeiterin im Bereich Seminarwesen und Publikationswesen neu eingestellt, zwei Kolleginnen befinden sich im Erziehungsurlaub, ein Mitarbeiter ist im wohlverdienten Ruhestand. Zwei Stellen wurden in 2021 nachbesetzt, bei zwei weiteren liefen befristete Arbeitsverhältnisse aus.

Seminare und Workshops

Die anhaltende Covid-19 Pandemie beeinflusste auch im Jahr 2021 das Seminarwesen stark. Dank eines Hygienekonzepts musste das Tagungszentrum nicht schließen und stand im ganzen Jahr für Seminare, Gruppentreffen und Forschungsaufenthalte zur Verfügung. Da viele Gäste nicht uneingeschränkt reisen konnte oder durften, entfielen aber gerade in der ersten Hälfte des Jahres viele Veranstaltungen.

Nach den ersten Versuchen in 2020 bot Schloss Dagstuhl wärend der Pandemie allen Seminaren die Option an, diese hybrid – also mit Teilnehmern, die vor Ort sind oder über Video-/Audioübertragung zugeschaltet sind – oder ganz virtuell als reine Videokonferenz durchzuführen. Insgesamt fanden dann 47 Seminare statt, davon allerdings nur 2 vollkommen vor Ort, 13 ganz online und die restlichen 32 hybrid. 21 Organisatorenteams haben ihre Seminare abgesagt, um sie später unter besseren Bedingungen in der typischen Dagstuhl Atmosphäre vor Ort abzuhalten. 9 Seminare wurden von Anfang des Jahres auf spätere Termine verschoben, die meist durch andere Absagen frei geworden waren. Die beiden geplanten Dagstuhl-Perspektiven-Workshops wurden nach 2022 verschoben – bei diesem Typ macht nach Meinung von Schloss Dagstuhl und den beteiligten Organisatoren nur eine Veranstaltung vor Ort mit einem breit gefächertem internationalem Publikum Sinn. Schließlich wurde auch ein Dagstuhl-Seminar nach 2022 verschoben, da durch Änderung einer Verordnung nicht klar war, welche pandemiebedingten Auflagen gelten würden. Immerhin konnten damit etwa 65 % der ursprünglich geplanten Veranstaltungen in irgendeiner Form durchgeführt werden. Da durch die Online-Angebote Seminare auch über der üblichen Größe stattfinden konnten, war der Rückgang auf etwa 75 % der durchschnittlichen Teilnehmer nicht so dramatisch wie der Rückgang der Veranstaltungen.

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

training, the “Dagstuhler Gespräche” (Dagstuhl Talks) series, and parts of the cooperation with the Heidelberg Laureate Forum had to be suspended again for 2021. These events are to be fully resumed as soon as the situation permits.

The Team

By the end of 2021, Schloss Dagstuhl had a total of 49 staff members corresponding to 39.28 full-time positions. In 2021, one trainee and one employee in the administration of Seminars and the publishing branch joined our team. Two employees took maternal leave or parental leave. One staff member took their well earned retirement. Two positions that became vacant were filled again, and the terms of two temporary positions ended.

Seminars and Workshops

In 2021, the still ongoing Covid-19 pandemic had a strong influence on the seminar department again. Due to a hygiene concept, the meeting center did not have to close down and was available for seminars, group meetings and research stays during the whole year. As many guests could not travel unrestrtictedly or where not allowed to travel, many events had to be canceled, especially in the first half of the year.

For that reason and after first attempts on 2020, during the pandemic, Schloss Dagstuhl offered all Seminars the option to be hybrid – i.e., with some participants that are on site at Schloss Dagstuhl and some that are included using video and audio transmission – or fully virtual as a video conference only. In fact, of the 47 seminars that took place, only two were on-site only, while 13 were remote only, and the remaining 32 were held in a hybrid fashion. 21 teams of organizers canceled their seminars in order to have them take place at a later point in time when conditions are improved and they can have the typical Dagstuhl atmosphere with everyone on site. 9 Seminars where moved from the beginning of the year to later dates which became available due to cancellations of other events. Both of the planned Dagstuhl Perspectives Workshops were moved to 2022 – Schloss Dagstuhl and the organizers felt that this type of event only makes sense if it has a wide variety of international participants on site. Finally, one Dagstuhl Seminar was moved to 2022 as well, as changes in local ordinances made it unclear exactly what precautionary measures would have had to be imposed on it due to the pandemic situation. At least, 65 % of the events that had been planned could take place in some form or another. As virtual events could accommodate a larger group of participants than usual, the number of guests overall was only reduced to 75 %, which was thus less dramatic than the reduction in the number of events.

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

Nationale Forschungsdateninfrastruktur für Datenwissenschaften und KI

Seit Oktober 2021 ist Schloss Dagstuhl mit den digitalen Infrastrukturen *Dagstuhl Publishing* und *dblp* Teil des Konsortiums NFDI4DataScience[2]. Ziel dieses Konsortiums mit über 15 Partnerinstitutionen ist es, für die Forschungsgebiete Data Science und Künstliche Intelligenz alle Schritte des komplexen und interdisziplinären Lebenszyklus der Forschungsdaten vollumfänglich zu unterstützen. Dies betrifft unter anderem die Schritte der Datenpublikation, der Erstellung von Metadaten und persistenten IDs, der Indexierung, sowie die Attributierung der Urheberschaft von Forschungsdaten.

Als Teil der im Aufbau befindlichen Nationalen Forschungsdateninfrastruktur (NFDI)[3] leistet Schloss Dagstuhl damit einen Beitrag, die wertvollen Datenbestände und Artefakte für das gesamte deutsche Wissenschaftssystem systematisch zu erschließen, nachhaltig zu sichern, zu vernetzen und für die Forschungsgemeinschaft nachnutzbar zu machen. Das Konsortium wird zunächst für 5 Jahre von Bund und Ländern finanziert.

Neue Rekordzahlen der Bibliographiedatenbank dblp

Das Jahr 2021 erwies sich erneut als das produktivste Jahr in der Geschichte der Datenbank. Die Arbeit des dblp-Teams wurde nahezu vollständig im Homeoffice verrichtet. So wurden innerhalb von zwölf Monaten mehr als 520 000 neue Publikationen indexiert. Die Neuaufnahmequote übertrifft damit zum fünften Mal in Folge die Rekordzahl aus dem vorausgangenen Jahr.

Auch bei der Anzahl manueller Korrekturen an den bestehenden Bibliographien konnte in 2021 ein neuer Rekord erzielt werden. Bei gleichbleibender Team-Größe wurden im Laufe des Jahres mehr als 72 000 Fehlerfälle bearbeitet. Dies entspricht einer Steigerung um 18,9% gegenüber dem Vorjahr.

Inzwischen registrieren die dblp-Webseiten im Schnitt mehr als 1,33 Millionen Seitenzugriffe pro Tag. Jeden Monat nutzen mehr als 750 000 verschiedene Nutzer aus aller Welt die dblp-Webservices.

Mehr Details zu dblp finden sich in Kapitel 3.

Dagstuhl Publishing

Wie in den Vorjahren haben die Open-Access-Publikationsaktivitäten auch in 2021 starken Zuspruch bekommen. So wurden in den Konferenzbandreihen LIPIcs und OASIcs zusammen 1 494 Publikationen innerhalb eines Jahres veröffentlicht. Zudem gab es auch in 2021 wieder viele Bewerbungen von wissenschaftlichen Konferenzen zur Veröffentlichung des Konferenzbandes in den Serien LIPIcs und OASIcs. Im 2021 konnte mit der Veröffentlichung von knapp 400 retro-digitalisierten Seminarberichten aus den Jahren 1990 bis 2003 auch ein langfristigstes Projekt

National Research Data Infrastructure for Data Science and AI

Since October 2021, Schloss Dagstuhl with its digital infrastructures *Dagstuhl Publishing* and *dblp* is a founding member of the consortium NFDI4DataScience[2]. With its more than 15 partner institutions, the goal of this consortium is to fully support all steps of the complex and interdisciplinary lifecycle of research data within the data science and artificial intelligence disciplines. This includes support for publishing the data, the creation of metadata and persistent IDs, indexing, and the attribution of its contributors.

As part of the National Research Data Infrastructure (NFDI)[3], which is currently under construction, Schloss Dagstuhl is therby contributing to foster findability, accessabilty, interoperability, and reuse of research data and artifacts for the German computer science community. The consortium is funded by the federal and state governments for an initial period of 5 years.

New Records of the dblp computer science bibliography

While for 2021 the work of the dblp team was carried out almost entirely in home office, the year once again proved to be the most productive year in the history of dblp. Within 12 months, more than 520,000 new publications were indexed. This is the fifth year in a row that the rate of new additions has exceeded the record number from the previous year.

Another record was set for the number of manual data corrections that have been made to existing author bibliographies. While the size of the team remained constant, more than 72,000 error cases have been handled and resolved during 2021. This represents an increase of 18.9% compared with the previous year.

Meanwhile, dblp websites receive on average more than 1.33 million page views per day. Eeach month, more than 750,000 distinct users from all over the world use the dblp web services.

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 2021. So in the conference proceedings series LIPIcs and OASIcs together, 1,494 publications were published within one year. Furthermore, LIPIcs again received and accepted proposals from several major scientific conferences. In 2021, the publication of almost 400 retro-digitised seminar reports from the years 1990 to 2003 also marked the completion of a long-term project. In addition, the sub-

[2] https://www.nfdi4datascience.de/
[3] https://www.nfdi.de/

abgeschlossen werden. Zudem wurde das 2019 eingeführte Einreichungssystem weiterentwickelt, welches die Arbeiten für Herausgeber, Editoren aber auch das Verlagsteam deutlich vereinfacht.

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

Zusammenarbeit mit dem Heidelberg Laureate Forum

Im Jahr 2021 fand das Heidelberg Laureate Forum[4] (HLF) nur virtuell statt. Dabei war Schloss Dagstuhl durch eine Informationsseite und ein Livechat in einer virtuellen Umgebung vertreten. Eine Wiederaufnahme der vollen Kooperation mit der Möglichkeit für ausgewählte HLF Teilnehmer, in der Woche vor oder nach dem HLF an Dagstuhl Seminaren teilnehmen zu können, ist geplant, sobald die Pandemielage das wieder zulässt.

Spender und Förderer

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

2021 erhielt die Bibliothek von mehreren Verlagshäusern erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 598 Bände als Spende, darunter 582 Monographien des Springer-Verlags im Wert von 46 627,76 €.

Baumaßnahmen und Renovierung

Der Vortragsraum „Saarbrücken“ erhielt einen ein neuen pflegeleichten Boden inklusive Bodensteckdosen, seine Decke wurde mit Dämmplatten ausgestattet, seine Beleuchtung modernisiert und eine zusätzliche Sitzbank geschaffen.

In den Betriebsferien wurden notwendige Malerarbeiten im ganzen Institut durchgeführt.

Ausstattung

Die bereits in 2020 begonnene Installation der Technik, um Seminare in einem hybriden Format, also mit einer Mischung aus Teilnehmenden vor Ort und per Videokonferenztechnologie Teilnehmenden, uneingeschränkt zu ermöglichen, wurde für die drei größten Hörsäle abgeschlossen. Sie sind nun mit Deckenmikrophonen und Kameras ausgestattet, die wahlweise den Raum oder den Vortragenden beziehungsweise einzelne Tafel- oder Whiteboardabschnitte erfassen können. Auch erhielten sie jeweils eine zusätzliche Anzeigemöglichkeit für eine Videokonferenz. Außerdem wurde der Vortragsraum „Saarbrücken“ mit einem neuen barrierearmen Rednerpult und einem großen Monitor ausgestattet.

mission system introduced in 2019 was further developed, which significantly simplifies the work for editors and the publishing team.

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

Joint Outreach with the Heidelberg Laureate Forum

In 2021, the Heidelberg Laureate Forum[4] (HLF) took place in an online format only. Schloss Dagstuhl was represented there with a website and a livechat in a virtual environment. As soon as the pandemic situation allows for it, the plan is to return to the full cooperation that offers selected HLF participants the chance to participate in a Dagstuhl Seminar the week before or after the HLF.

Sponsors and Donors

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

The center's research library received a large number of book donations from several publishing houses. The number of donated volumes totaled 598, including 582 monographs at a total value of 46,627.76 € donated by Springer Science+Business Media publishing house.

Construction Work and Renovation

Lecture hall "Saarbrücken" got a new low-maintenance flooring with floor power outlets, sound insulation panels on the ceiling, modernized lighting, and an additional seat bench.

During the vacation close-down, necessary paintwork all over the institute was performed.

Facilities

The installation of technology to facilitate seminars in a hybrid format unrestrictedly, i.e., with a mixture of on-site participants and remote participants using videoconferencing, which already started in 2020, has been completed for the three biggest lecture halls. They are now equipped with ceiling microphones and cameras that pick up the room or the speaker or specific parts of the white- or blackboards, respectively. They were also equipped with an additional projector or monitor each to display a video conference. The lecture hall "Saarbrücken" was also equipped with a new low-barrier lectern and a big monitor.

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

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 2021 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 the exchange of research views and findings in a specific cutting-edge field of informatics between promising young scientists and the international elite of the research area.

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 past guests are 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 have 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 from the Dagstuhl website.[5]

Chapter 6 contains a collection of the summaries of the 2021 Seminars and Perspectives Workshops. Chapter 14 provides a comprehensive list of all events that took place during 2021, and a seminar program covering the upcoming 24 months is available on the Dagstuhl website.

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

2.2 Dagstuhl-Perspektiven-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.

Aufgrund der Covid-19 Pandemie konnten die beiden für 2021 geplanten Dagstuhl-Perspektiven Workshops leider nicht stattfinden. Sie wurden aber ins Jahr 2022 verschoben.

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]

2.2 *Dagstuhl Perspectives Workshops*

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, and
- trigger innovation processes.

Due to the Covid-19 pandemic, the two Dagstuhl Perspectives Workshops scheduled for 2021 could not take place. They have been rescheduled to take place in 2022.

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 Workshops can be found on our website.[7]

2.3 Einreichung der Anträge und Begutachtungsverfahren

Die gleichbleibend hohe Qualität der Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops wird durch Auswahl der Anträge gewährleistet, die aus Sicht von Schloss Dagstuhl das größte Potential haben, abseits etablierter Konferenzen neue und wichtige Forschungsprobleme mit Wissenschaftlern aus oft unterschiedlichen Gebieten zu identifizieren und zeitgleich mögliche Methoden und Lösungsansätze zu diskutieren.

Das Zentrum erbittet zweimal im Jahr Themenvor-

2.3 *Proposal Submission and Review Process*

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

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

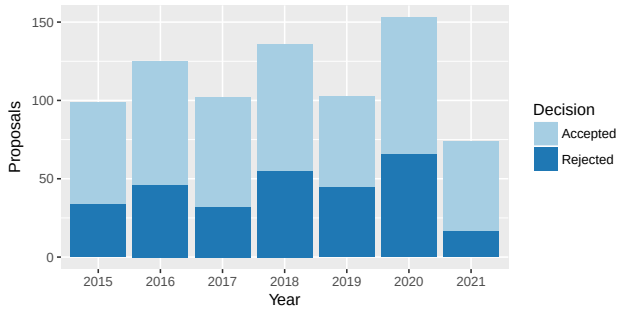

Fig. 2.1
Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2015–2021.

schläge 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.

In 2021 fanden wie gewohnt zwei Antragsrunden zu den üblichen Terminen statt. Die erste Antragsrunde war jedoch auf die Neubeantragung abgesagter Seminare aus 2020 beschränkt. Die in dieser Runde eingereichten 22 Anträge wurden alle angenommen. Für die Begutachtung im Juni wurden mit 52 Anträge ähnlich viele wie sonst auch eingereicht. Dabei waren 12 Anträge Neubeantragungen für Seminare, die durch die Pandemie nicht oder nur erheblich eingeschränkt stattfinden konnten.

Damit sind die statistischen Daten über die Anträge nicht direkt mit denen aus dem Vorjahr vergleichbar (siehe Fig. 2.1).

Unter den 57 in 2021 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 sind 40 Seminare (hier und im Folgenden wird, sofern nicht anders angegeben, das Wort „Seminar“ sowohl für Dagstuhl-Seminare als auch für Dagstuhl-Perspektiven-Workshops verwendet) für 2022 eingeplant; 17 konnte erst 2023 geplant werden. Diese langen Vorlaufszeiten entstanden maßgeblich durch die Verschiebung und Absage von Seminaren in 2020 und 2021.

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.

As ususal, there were two proposal rounds in 2021 which both took place on the customary dates. However, the first proposal round was limited to proposals that proposed anew seminars which originally had been accepted for 2020 but had to be cancelled. All 22 proposals handed in in that round were accepted. For review in June, 52 proposals were handed in, which is a rather normal number. Of these, 12 were proposals proposing a seminar anew which due to the pandemic had only taken place in a rather limited fashion or not taken place at all.

Thus, the statistical data on proposals is not directly comparable to data from the preceding year (see Fig. 2.1).

Among the 57 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2021 there is – as in the past years – a wide variation with regard to length and size (see Fig. 2.3). Of these seminars, 40 were scheduled to take place in 2022; 17 could only be scheduled for 2023 (here and in the following, the word "seminar" is meant to include both Dagstuhl Seminars and Dagstuhl Perspektives Workshops, if not specified otherwise). These long lead times are mostly due to postponement or cancellation of Seminars in 2020 and 2021.

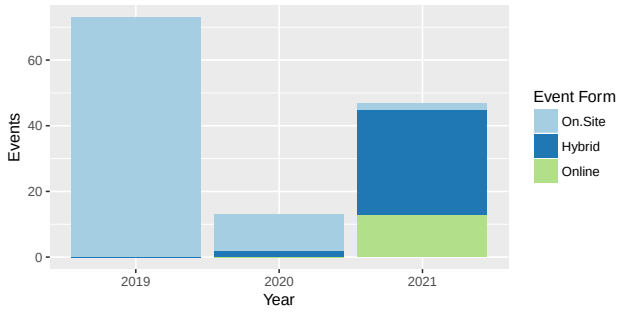

Fig. 2.2
Number of seminars in 2019–2021by form.

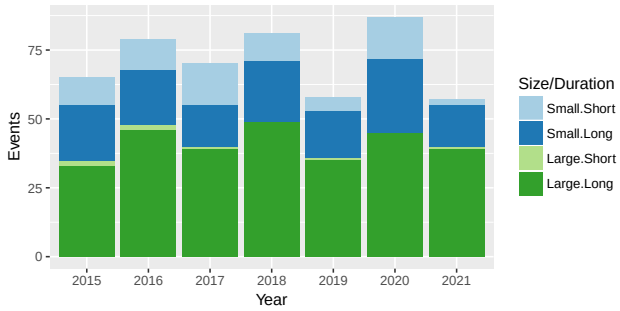

Fig. 2.3
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2015–2021.
Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

Seminar-Programm 2021

2.4 *The Seminar Program in 2021*

Hier und im Folgenden ist zu beachten, dass durch die pandemiebedingte Reduzierung der Aktivitäten des Tagungszentrums die Zahlen nicht mit den Zahlen der Vorjahre vergleichbar sind.

Grundsätzlich kann Schloss Dagstuhl in jeder Woche zwei Seminare mit insgesamt etwa 75 Teilnehmern beherbergen.

In 32 von 48 Wochen, in denen das Tagungszentrum 2021 geöffnet war, fand mindestens ein Dagstuhl-Seminar statt. In 15 Wochen waren es sogar zwei. In 6 Wochen war das Zentrum nur durch andere Veranstaltungen belegt. In 10 Wochen fand keine Veranstaltung statt.

2021 fanden insgesamt 47 Dagstuhl-Seminare statt. In Fig. 2.5 ist die Entwicklung der vergangenen Jahre dargestellt. Für die Zeit der Pandemie bietet Schloss Dagstuhl allen Seminarorganisatoren die Option an, diese hybrid – also mit Teilnehmern, die vor Ort sind oder über Video-/Audioübertragung zugeschaltet sind – oder ganz virtuell als reine Videokonferenz durchzuführen. Fig. 2.2 gibt einen Überblick, wie viele Veranstaltungen on-site, hybrid oder online stattgefunden haben.

It should be noted here and in the following that due to the pandemic-induced reduction in the activities of the meeting center, the figures are not comparable with those of previous years.

In principle, Schloss Dagstuhl can host two seminars each week with a total of about 75 participants.

At least one Dagstuhl Seminar was held in 32 of the 48 weeks the center was open in 2021. In 15 of those weeks, there were in fact two seminars in parallel. In 6 weeks, there were exclusively other events scheduled. In the remaining 10 weeks, no events took place.

Altogether, there were 47 Dagstuhl Seminars in 2021. Fig. 2.5 shows the evolution in recent years.

For the duration of the pandemic, Schloss Dagstuhl offers all Seminars the option to be hybrid – i.e., with some participants that are on site at Schloss Dagstuhl and some that are included using video and audio transmission – or fully virtual as a video conference only. Fig. 2.2 provides an overview of how many events took place on site, in a hybrid setting or online, respectively.

Angaben zu Teilnehmern und Organisatoren

2.5 *Participant and Organizer Data*

Bei den folgenden Vergleichen mit den Vorjahren ignorieren wir das Vorjahr 2020, da pandemiebedingt die Anzahl der Gäste sehr niedrig war und damit Aufgrund der

It should be noted here that, due to the pandemic, the number of guests was very low in 2020, so that due to the small sample size, all statistical data from 2020 are of

geringen Stichprobengröße alle statistischen Angaben dieses Jahrs von geringer Aussagekraft waren. Die absoluten Zahlen des Jahrs 2021 weichen aber immer noch erheblich von den Zahlen der Jahre 2015 bis 2019 ab.

Insgesamt nahmen 2293 Wissenschaftler an Veranstaltungen teil, davon 1894 an Dagstuhl-Seminaren. Jedoch waren nur 731 Wissenschaftler – davon 529 bei Seminaren – vor Ort. Grundsätzlich werden hier alle Teilnehmer gezählt, unabhängig davon, ob sie vor Ort oder online teilgenommen haben. Einen Eindruck, wieviele Teilnehmer vor Ort waren bzw. remote teilgenommen haben, vermittelt Fig. 2.4.

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 betrachteten Vorjahren, in 2021 wieder fast die Hälfte (49 %, 894 von 1 827) der Wissenschaftler das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während weitere 17 % der Wissenschaftler an nur einem Seminar in den Jahren vorher teilgenommen hatten, weitere 9 % nur an zweien. Betrachtet man ausschließlich Teilnehmer, die vor Ort angereist waren, so erhält man bis auf die erstmaligen Teilnahmen, die auf 44 % sinken, ein vergleichbares Ergebnis.

Ein wenig andere Zahlen leiten sich aus unserer Gastumfrage ab. Hier ergibt sich, dass etwa 43 % der Antwortenden 2021 das erste Mal, 15 % zum zweiten Mal und weitere 10 % zum dritten Mal (siehe Fig. 2.6a) teilgenommen haben.

Ein beträchtlicher Anteil der Gäste besteht aus jungen Wissenschaftlern, die am Anfang ihrer Karriere stehen, und für die der Aufenthalt in Dagstuhl oftmals prägend ist für den weiteren Verlauf ihres Lebenswegs. Etwa 35 % der Gäste der Dagstuhl-Seminare in 2021, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein (siehe Fig. 2.6b). Wie in allen Vorjahren – bis auf das letzte – hatte Schloss Dagstuhl mit etwa einem Drittel eine ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern, die im Laufe der Jahre (mit Ausnahme des letzten) relativ konstant geblieben war, 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 2021 wieder auf einem gewohnt hohen Niveau. Das Diagramm in Fig. 2.6c zeigt die regionale Verteilung der Gäste für 2021 bei Dagstuhl-Seminaren. Mehr Details können Kapitel 13 entnommen werden.

In 2021 waren etwa 92 % aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt und rund 29 % aller Organisatoren waren Frauen (siehe Fig. 2.7a). Der Anteil an weiblichen Seminarteilnehmern war mit 24 % wieder auf dem Niveau von 2019 und damit höher als in allen anderen betrachteten Jahren (siehe Fig. 2.7b).

little significance, and will thus be ignored in the following. However, the numbers for 2021 still deviate significantly from the numbers for the years between 2015 and 2019.

Of the more than 2,293 scientists who participated in events, 1,894 participated in seminars. Only 731 scientists – of these 529 participating in seminars – participated on site. Basically, all participants are counted here, regardless of whether they participated on site or online. Fig. 2.4 illustrates how many participants were on site and how many participated remotely.

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 2021, almost half (894 of 1,827, or 49 %) of the researchers were first-time visitors to Dagstuhl. About an additional 17 % of the participating researchers had already attended one previous seminar in the years before, and another 9 % had already attended two. Looking only at on-site participants, similar numbers are obtained except for the first-time visits, which drop to 44 %.

Slightly different numbers are obtained from our guest survey: About 43 % of the responders were first-time visitors, an additional 15 % state their second visit, and yet another 10 % their third (see Figure 2.6a).

A substantial number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of lifelong value. Approximately 35 % of 2021 Dagstuhl Seminar survey respondents self-classified as junior (see Fig. 2.6b). Like in the previous years – except 2020 – Schloss Dagstuhl had a balanced proportion of junior to senior researchers of about a third, 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 guests with a non-German affiliation in Dagstuhl Seminars was back to it's usual high level. The chart in Fig. 2.6c shows the regional distribution of our Dagstuhl Seminar guests in 2021. For a detailed breakdown please refer to Chapter 13.

In 2021, about 92 % of all organizer teams in our scientific seminar program were mixed with respect to gender and about 29 % of all organizers were women (see Fig. 2.7a). The percentage of female seminar participants was back to the level of 2019, at 24 %, and thus higher than in all the other considered years (see Fig. 2.7b).

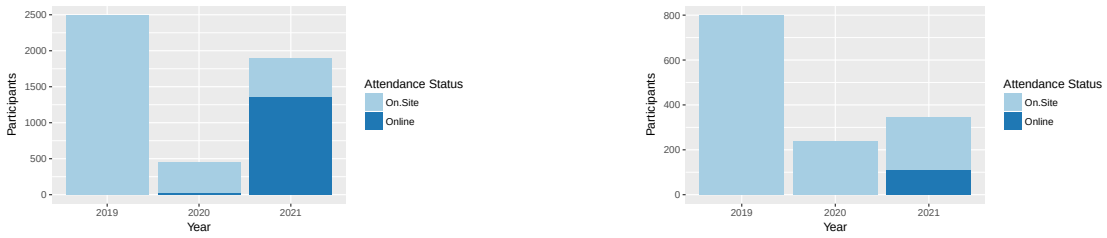

(a) Distribution of participants of seminars in group A.

(b) Distribution of participants of seminars in group B.

Fig. 2.4
Number of participants by attendance status and group. Group A = Dagstuhl Seminars and Dagstuhl-Perspectives-Workshops, group B = all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).

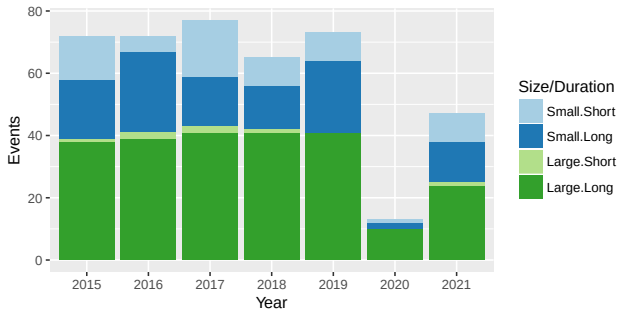

Fig. 2.5
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2015–2021.
Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

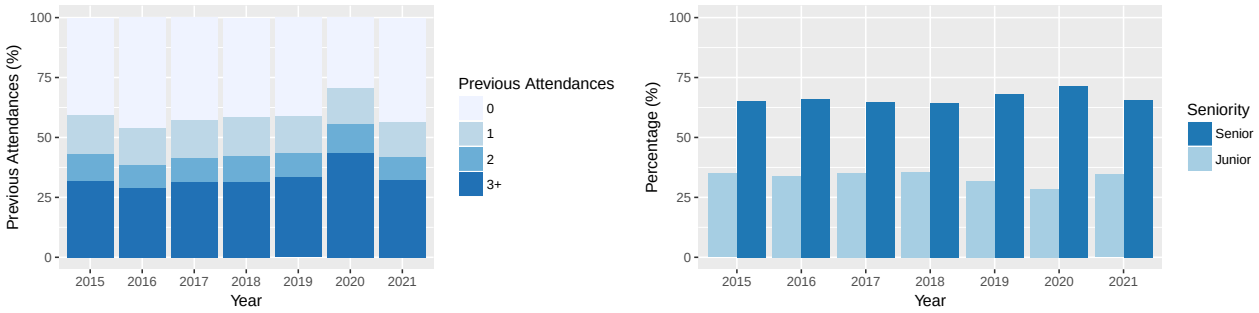

(a) Distribution of the number of previous attendances of participants, according to survey data.

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

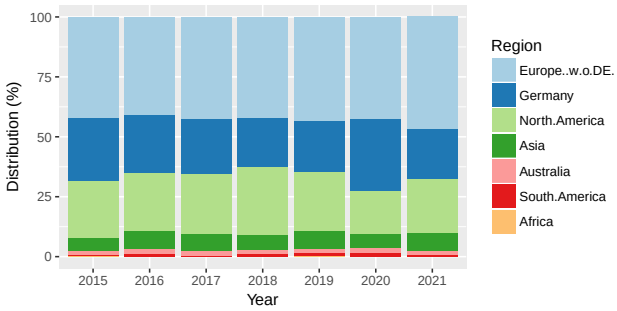

(c) Distribution of the origin of participants by region.

Fig. 2.6
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2015–2021.

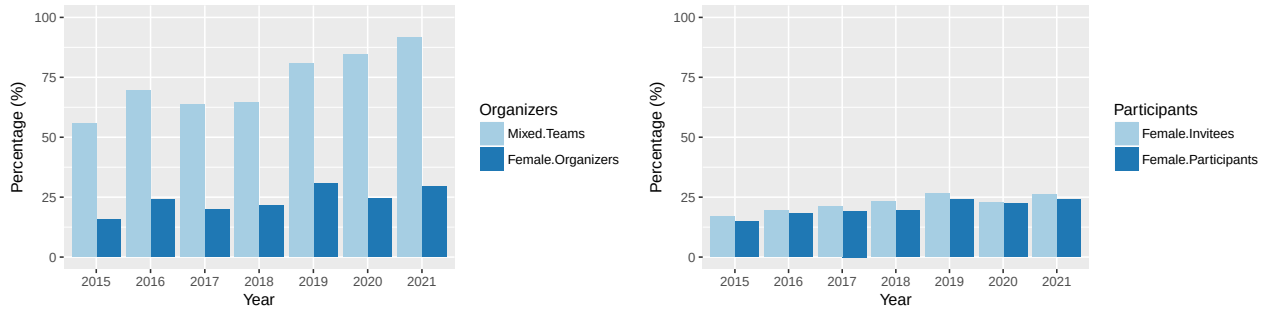

(a) Female organizers and mixed-gender organizer teams.

(b) Female invitees and participants.

Fig. 2.7

Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2015–2021.

Themen und Forschungsgebiete 2.6 *Topics and Research Areas*

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

Im Folgenden sind beispielhaft einige thematische Schwerpunkte und dazugehörige Seminare aufgeführt. Die Aufzählung der Themen und Seminare hat keinen Anspruch auf Vollständigkeit und ist lediglich ein Versuch, einen kurzen Einblick in das facettenreiche 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 2021.

Trotz des etwa auf 65 % reduzierten Seminarporgramms deckten die Themen große Teile der Informatik ab. Von den theoretischen Grundlagen, etwa *Computational Complexity of Discrete Problems* (21121), über mehr angewandte Themen wie *Secure Compilation* (21481) bis zu praktisch relevanten Themen wie *Managing Industrial Control Systems Security Risks for Cyber Insurance* (21451) war wieder alles vertreten. Anwendungen vom omnipräsenten Thema Machine Learning waren auch Bestandteil einiger Seminare wie zum Beispiel *Machine Learning in Sports* (21411). Auch gesellschaftspolitische relevante Fragen werden in Dagstuhl diskutiert wie die Seminare *Towards Climate-Friendly Internet Research* (21272), in dem auch um das Vermeiden unnötiger Reisen ging, und *Digital Disinformation: Taxonomy, Impact, Mitigation, and Regulation* (21402), in dem es auch um Desinformation ging, zeigen.

Leider sind 2021 auch wieder viele spannende Themen der Pandemie zum Opfer gefallen, etwa *Media Forensics and the Challenge of Big Data*, *Privacy in Speech and Language Technology* oder *Algorithms and Law*. Wir sind aber zuversichtlich, dass diese in naher Zukunft unser Programm erneut bereichern werden.

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 2021. Neither the list of focal points nor the list of seminars is exhaustive. It merely attempts to offer a brief insight into the multifarious scientific seminar program of 2021. Chapter 6, with the summary of the Seminars and Perspectives Workshops, provides a full overview of the 2021 scientific seminar program.

Despite the seminar program being reduced to about 65 %, the topics covered large parts of computer science. From the theoretical basics, such as *Computational Complexity of Discrete Problems* (21121), to more applied topics like *Secure Compilation* (21481) to practically relevant topics like *Managing Industrial Control Systems Security Risks for Cyber Insurance* (21451), everything was represented again. Applications of the omnipresent topic of machine learning were also a part of some seminars like for example *Machine Learning in Sports* (21411). Questions of societal relevance are being discussed in Dagstuhl as well, as demonstrated by the seminars *Towards Climate-Friendly Internet Research* (21272), which also looked at avoiding unnecessary travel, and *Digital Disinformation: Taxonomy, Impact, Mitigation, and Regulation* (21402), which also looked at disinformation.

Unfortunately, many exciting topics fell victim to the pandemic in 2021 again, for example, *Media Forensics and the Challenge of Big Data*, *Privacy in Speech and Language Technology*, or *Algorithms and Law*. However, we are confident that they will enrich our program again in the near future.

This brief selection of seminars should not draw attention from the fact that each of 2021's seminars addressed important topics which were discussed by the involved researchers with great commitment and hence pushed

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der in 2021 veranstalteten Seminare wichtige Themen adressiert hat, die von den beteiligten Wissenschaftlern mit großem Engagement diskutiert wurden und so die weitere Entwicklung in den einzelnen Gebieten wieder ein gutes Stück weitergebracht hat. An dieser Stelle möchten wir auch allen Organisatoren danken, die auch unter den widrigen Umständen der Covid-19 Pandemie Dagstuhl-Seminare geleitet haben – auch wenn zu großen Teilen das Erlebnis einer repräsentativen, internationalen Gruppe von Wissenschaftlern vor Ort nicht möglich war.

forward the development in the individual areas. We would like to use this opportunity to thank all organizers who braved the adverse circumstances of the Covid-19 pandemic to lead Dagstuhl Seminars – even if large parts of the experience of having a representative group of international scientists on site were not possible.

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

Further Event Types

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.
- 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. In den zur Verfügung gestellten PDF-Dokumente werden die statistischen Ergebnisse mit Hilfe von aussagekräftigen Diagrammen aufbereitet.

Fig. 2.8 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2021 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 660 Fragebögen, welche

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 the pdf files with the results, the statistics are visualized using illuminative diagrams.

Fig. 2.8 shows the satisfaction of responding participants in 2021 with regard to selected aspects of their stay. The results were compiled from 660 questionnaires, representing the responses of about 35 % of all 1,898 participants. These excellent results are not only a recognition

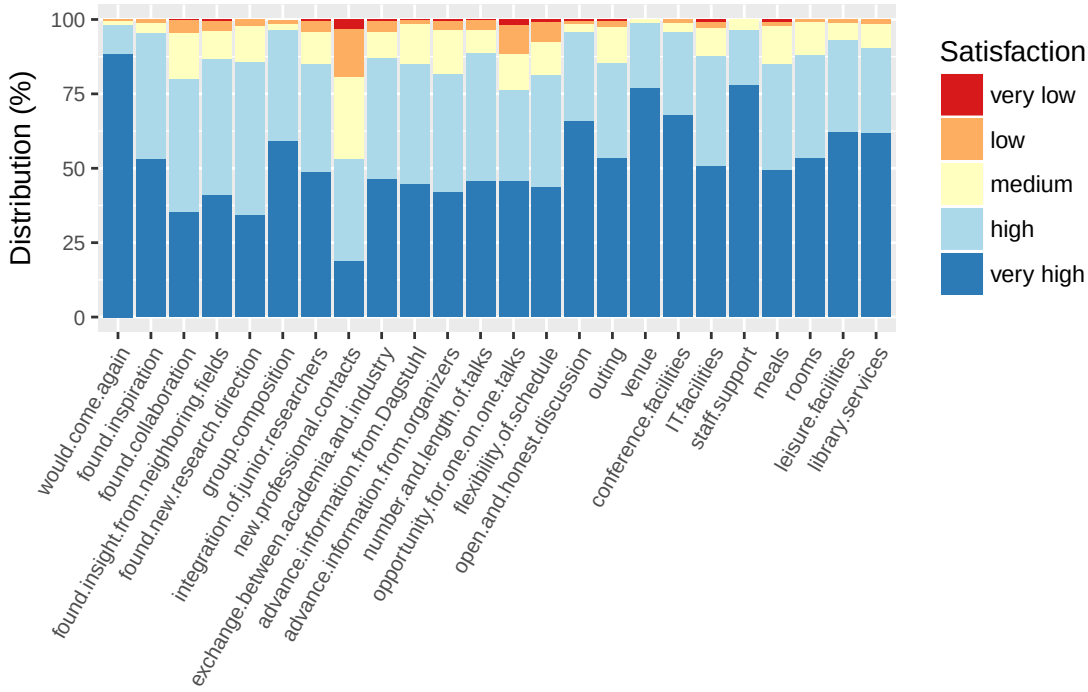

Fig. 2.8
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2021. According to survey results.

die Meinung von etwa 35 % der 1 898 Teilnehmer repräsentieren. Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich. Die Rücklaufrate der Fragebögen war damit ungewöhnlich niedrig. Es kann vermutet werden, dass es an dem hohen Anteil der Teilnehmer liegt, die remote teilgenommen haben.

Als Teil des Einladungsprozesses der Dagstuhl-Seminare und der Dagstuhl-Perspektiven-Workshops wird die Liste der von den Organisatoren zur Einladung Vorgeschlagenen von Schloss Dagstuhl auf eine ausgewogene Zusammensetzung geprüft, bevor Schloss Dagstuhl Einladungen ausspricht. Mittels einer täglich aktualisierten Webseite bietet Schloss Dagstuhl allen Organisatoren einen direkten Einblick in den Status der eingeladenen Gäste bezüglich Zu- oder Absage.

of the center's past work but also pose a challenge to its future work. The turn-in rate of the questionaires was thus unusually low. It can be suspected that this is due to the high ratio of remote participants.

During the invitation process for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, the Organizers compile a list of proposed invitees which is reviewed by Schloss Dagstuhl to check it for a balanced composition before Schloss Dagstuhl extends invitations. Via a dedicated webpage that is updated daily, Schloss Dagstuhl gives the organizers direct access to view the status of invitee replies.

Auslastung des Zentrums 2.9 *Utilization of the Center*

2021 war die Auslastung aufgrund der Covid-19 Pandemie zwar höher als im Vorjahr aber immer noch geringer als in den Jahren davor. Es gab 2021 insgesamt 3 123 Gasttage, wobei 2 397 Gasttage auf Dagstuhl-Seminare entfielen. Im Gegensatz zu den meisten anderen Statistiken zählen hier nur die Tage, an denen die Gäste tatsächlich vor Ort waren. Es fanden im Berichtsjahr 63 Veranstaltungen statt, bei denen wenigstens ein paar Teilnehmer vor Ort waren. Insgesamt gab es 762 Gäste vor Ort. Weitere Details können Kapitel 13 entnommen werden.

Die Wochenenden blieben 2021 ebenso unbelegt wie zwei Wochen zu Weihnachten und 2 Wochen im Sommer, welche zu Instandhaltungs- und Verwaltungsarbeiten

Due to the Covid-19 pandemic, the capacity utilization in 2021 was higher than in the previous year, but still lower than in the years before 2020. There were 3,123 overnight stays in total in 2021, with 2,397 overnight stays in Dagstuhl Seminars. As oposed to in most of the other statistics, only days in which the guests where actually on site are counted here. The center hosted a total of 63 events where at least a few people where on site in 2021. In total, there were 762 guests on site. See Chapter 13 for further details.

Weekends were kept free in 2021, as well as a week at the beginning of the year and a week at the end of the year, this time being required for maintenance work to building facilities and administrative work. Of the 48 weeks that

benötigt wurden. Aus den 48 prinzipiell zur Verfügung stehenden Wochen waren in 32 Wochen tatsächlich Gäste vor Ort.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2021 einschließlich Dagstuhl-Seminaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-Seminaren und Veranstaltungen (z.B. Sommerschulen), bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 14. Auf unserer Webseite ist ein Kalender[8] verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

where available in priciple, 32 weeks saw guests actually on site.

A comprehensive listing of all events at Schloss Dagstuhl in 2021, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and host-only events such as meetings and summer schools can be found in Chapter 14. See the Schloss Dagstuhl website to view our calendar[8] of upcoming events and further information and materials on all events past, present and future.

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

3 Bibliographiedatenbank dblp
dblp computer science bibliography

Offene Bibliographiedaten für die Informatik

3.1

Open Bibliographic Data in Computer Science

Moderne Informatik-Forschung benötigt den unmittelbaren und umfassenden Zugriff auf aktuelle Publikationen, um den Bedürfnissen in einer sich immer schneller entwickelnden und immer komplexer werdenden Forschungslandschaft gerecht zu werden. Doch nicht nur im Alltag der Forschenden, 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, 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 Forschenden und Institutionen bei der Berichtspflicht durch die Sammlung und Aufbereitung von qualitätsgesicherten Publikationslisten
- Unterstützung von Forschungsfördernden und Entscheidungstragenden durch das öffentliche Verfügbarmachen von nach Daten-Facetten aufgeschlüsselten Publikationsnachweisen

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

Modern computer science research requires 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.[9] 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 46 000 Treffer, Semantic Scholar liefert 2 430; im Einzelnen weisen SpringerLink ca. 5 200 Artikel, Elsevier ScienceDirect über 1 150 Artikel, die ACM Digital Library ca. 4 800 Artikel und IEEE Xplore über 3 100 Artikel nach.
The search term „dblp“ results in more than 46 000 hits at Google Scholar and 2 430 hits at Semantic Scholar; in particular, SpringerLink lists about 5 200 articles, Elsevier ScienceDirect lists more than 1 150 articles, the ACM Digital Library lists 4 800 articles, and IEEE Xplore lists more than 3 100 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 dauerhafte Personalstellen im wissenschaftlichen Stab geschaffen, die hauptamtlich mit der 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 von Bund und Ländern weitere Mittel für den Betrieb von dblp bereitgestellt und eine neue, eigens eingerichtete Außenstelle von Schloss Dagstuhl auf dem Campus II der Universität Trier angesiedelt. Betrieb und Erforschung der Datenbank erfolgen dabei weiterhin in enger Kooperation mit dem Fach Informatikwissenschaften der Universität Trier.

Das dblp-Team besteht mitlerweile aus 8 Vollzeitäquivalenten, welche an der redaktionellen, technischen und wissenschaftlichen Verbesserung der Infrastruktur arbeiten. Das Team konnte 2021 mehr als 520 000 neue Publikationen indizieren (siehe Abschnitt 3.3). Gleichzeitig konnten eine neue Rekordzahl an Fehlern im Bestand korrigiert werden (siehe Abschnitt 3.4). Parallel hierzu wird kontinuierlich an neuen Funktionen gearbeitet.

The cooperation between Schloss Dagstuhl and the dblp computer science bibliography – originally developed at the University of Trier – has been existing since late 2010. The commitment of Schloss Dagstuhl to dblp, initially funded by a project of the Leibniz Competition, is beeing funded directly by Schloss Dagstuhl since June 2013. As part of the consolidation of this cooperation, permanent scientific staff positions – assigned fulltime to the support and development of dblp – were created. The dblp advisory board (cf. 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 at Schloss Dagstuhl took place just in time for dblp's 25th anniversary. At the same time, Dagstuhl's funding was increased to support the operation of dblp and a new Schloss Dagstuhl branch office for the dblp team was 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.

The dblp team – which had been a one-person project for more than a decade – now consists of 8 full-time equivalent staff members working on the editorial, technical, and scientific improvement of the infrastructure. In 2021, the team handled the indexing of more than 520,000 new publications (see Section 3.3) but also set a new record with respect to the curation of existing data (see Section 3.4). Parallel to the ongoing work with the data, services are continuously improved and new features are implemented.

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 Front Matter, Münster, 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, United States of America
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 2021.

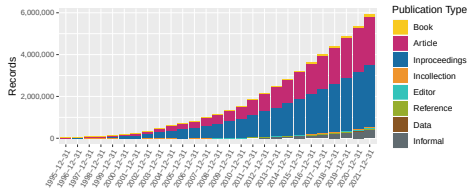

(a) Total number of records by year and type

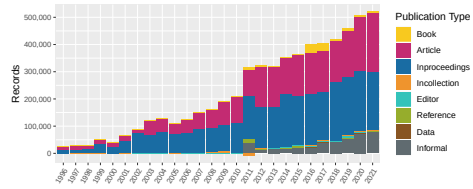

(b) New records by year and type

Fig. 3.2
Development of the dblp data stock.

Statistiken der Datenakquise 3.3 *Data Acquisition Statistics*

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

Obwohl die Arbeit des dblp-Teams nahezu vollständig im Homeoffice verrichtet wurde erwies sich 2021 erneut als das produktivste Jahr in der Geschichte der Datenbank. So wurden innerhalb von zwölf Monaten über 520 000 neue Publikationen indexiert. Das entspricht mehr als 2 000 neuen Publikationen pro Arbeitstag. Die Neuaufnahmequote übertrifft damit zum fünften Mal in Folge die Rekordzahl aus dem vorausgangenen Jahr. Die neu aufgenommenen Einträge verteilen sich zu 41,5% auf Konferenzbeiträge, zu 41,3% auf Journalartikel, zu 15,4% auf Preprints und „graue“ Literatur, sowie zu 1,9% auf andere Publikationstypen wie etwa Monographien und Dissertationen. Am Ende des Jahres waren über 5,9 Millionen Publikationen aus den verschiedenen Teilgebieten der Informatik indexiert.

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 executed almost exclusively by hand.

Although for 2021 the work of the dblp team has been carried out almost entirely in the home office, the past year once again proved to be the most productive year in the history of dblp. Within 12 months, more than 520,000 new publications were indexed. This figure corresponds to more than 2,000 new records for each working day of the year. Hence, for the fifth year in a row, the rate of new additions has exceeded the record number from the previous year. This year's new records consist of 41.5% conference papers, 41.3% journal articles, 15.4% preprints and "grey" literature, and 1.9% further publication types like monographs and PhD theses. By the end of 2021, more than 5.9 million publications were indexed by dblp.

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

Statistiken der Datenkuration 3.4 *Data Curation Statistics*

Ein Hauptziel der intensiven Datenpflege des dblp-Teams ist es sicherzustellen, dass die bereitgestellten Bibliographien so korrekt und vollständig wie möglich sind. Das bedeutet, dass alle Publikationen eines Forschen-

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

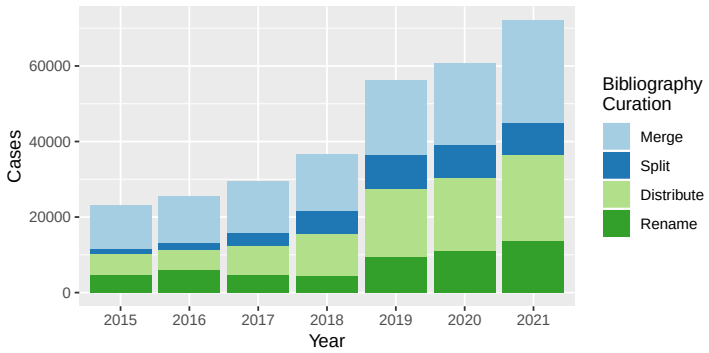

Fig. 3.3
Curation of existing dblp author bibliographies. The figures give the number of distinct edit cases (measured between the first and the last day of every given year) where a dblp team member manually corrected the assignment of publications within dblp author bibliographies. We distinguish between 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.

den in nur einer einzigen Bibliographie aufgeführt sein sollen und dass diese Bibliographie auch nur Publikationen des jeweiligen Forschenden listen soll. Dies zu gewährleisten kann ziemlich herausfordernd sein, und trotz der Bemühungen des Teams werden regelmäßig Publikationen einer falschen Bibliographie zugeordnet. Aus diesem Grund wird der dblp-Datenbestand kontinuierlich von der dblp-Redaktion überprüft und eventuelle Zuordnungs- oder Daten-Fehler korrigiert.

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.

Bei der Anzahl manueller durchgeführter Korrekturen an den bestehenden Bibliographien konnte in 2021 ein neuer Rekord erzielt werden. Trotz gleichbleibender Team-Größe wurden im Laufe des Jahres mehr als 72 000 Fehlerfälle bearbeitet. Dies entspricht einer Steigerung um 18,9% gegenüber dem Vorjahr. Fig. 3.3 zeigt die Anzahl der gelösten Fehlerfälle in den letzten Jahren. Hauptfaktoren für die kontinuerliche Steigerung sind verbesserte interne Werkzeuge zum Erkennen potentiellen Fehlern und eine verbesserte Verfügbarkeit von ORCID[10]-Daten. ORCIDs sind eindeutige Kennungen, die Autoren mit ihren Publikationen verbinden können. Durch den Vergleich von ORCID-Daten mit dblp-Bibliographien gelingt es, in großer Zahl Zuordnungsfehler aufzudecken.

that a bibliography should only list publications from that specific author. It can be quite challanging to ensure this level of quality, and despite the best efforts of the dblp team, we regularly assign publications to the wrong bibliography. Because of this, our editorial team constantly checks our data and corrects such defects.

While specialized algorithms help the team to uncover and identify the nature of defects 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.

In 2021, a new record was set for the number of manual data corrections that have been made to existing author bibliographies. While the size of the team remained constant, more than 72,000 error cases have been handled and resolved. This represents an increase of 18.9% compared with the previous year. Figure 3.3 shows the number of resolved defect cases during the past few years. The primary factors for this increase are more efficient internal tools that point out potential errors to us and an increased availability of ORCID[10] data. ORCID is a unique identifier that authors can attach to their publications. Comparing ORCID data with our bibliographies helps us to identify errors in large numbers.

[10] https://orcid.org

	Trier 1		Trier 2		Dagstuhl		Total		
	2020	2021	2020	2021	2020	2021	2020	2021	%
page views per day	712,228	644,837	40,523	18,206	486,166	672,274	1,238,918	1,335,318	+7.8
user sessions (visits) per day	31,450	20,134	414	297	36,849	31,537	68,715	51,969	-24.4
page views per user session	22.6	32.0	97.7	61.1	13.2	21.3	18.0	25.7	+42.5
distinct users (IPs) per month	473,014	257,769	4,433	3,512	578,492	492,268	1,055,941	753,550	-28.6
data served per month	2,243.6 GB	2,073.9 GB	170.3 GB	88.4 GB	1,402.1 GB	1,433.5 GB	3,816.0 GB	3,595.8 GB	-5.8

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 3.5 *Usage Statistics*

Im Jahr 2021 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. Mittlerweile ist Server Dagstuhl der mit Abstand am meisten genutze Server bezüglich der Anzahl der Besucher sowie der Platzierung bei Google.

Insgesamt konnten die Nutzungszahlen in 2021 gegenüber dem Vorjahr leicht gesteigert werden. Insbesondere die Zahl der Seitenaufrufe pro Tag erhöhte sich gegenüber dem Vorjahr um 7,8 Prozent auf nunmehr über 1,33 Millionen. Bei Bestimmung der unterschiedlichen Nutzenden pro Monat wurde 2021 die Berechnungsgrundlage geändert um zukünftig eine Doppel-Zählung von IP auf unterscheidlichen Servern zu vermeiden. Die Zahlen von 2020 und 2021 sind daher nur eingeschränkt miteinander vergleichbar. Demnach werden die dblp-Webservices insgesamt von mehr als 750 000 Nutzenden pro Monat besucht. 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 2021, 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 2020, server Dagstuhl overtook Trier 1 in respect to the number of visitors as well as the Google search ranking.

Overall, the total usage in 2021 slightly increased in 2021. In particular, the dblp websites receive on average more than 1.33 million page views per day. That is an increase by 7.8% compared with the previous year. In 2021, the counting method to determine distinct users has been changed in order to avoid double counting of the same IP across diferent servers. Figures for 2020 and 2021 are therefore only comparable with each other to a limited extent. In total, more than 750,000 distinct users from all over the world did use the dblp web services. 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*[11] 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 2021 wurde für 28 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*[12] 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 2021 wurde kein Buch in der Reihe veröffentlicht.

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*[11] 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 2021, 28 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*[12] 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 2021, no volume was published in the series.

[11] https://www.dagstuhl.de/dagrep
[12] https://www.dagstuhl.de/dfu

Dagstuhl Publishing

Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*[13] die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Dagstuhl-Perspektiven-Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 11.4) fungiert hier ebenfalls als Herausgebergremium. In 2021 wurde eine Ausgabe mit zwei Manifestos veröffentlicht (siehe Fig. 4.1).

DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*[14] werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2021 wurde die siebte Ausgabe mit zwei Heften und insgesamt 18 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. Im Jahr 2015 gab es zum Beispiel einen Perspektiven-Workshop mit dem Titel „Artifact Evaluation for Publications“[15], der 2016 durch zwei Seminare ergänzt wurde: „Reproducibility of Data-Oriented Experiments in e-Science“[16] und „Rethinking Experimental Methods in Computing“[17].

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*[13] in an Open Access manner. The Scientific Directorate (see Fig. 11.4) acts as the editorial board of the journal. In 2021 one volume with two manifestos was published (see Fig. 4.1).

DARTS: Dagstuhl Artifacts Series

The *DARTS* series[14] publishes evaluated research data and artifacts. It is organized as a periodical. In 2021, the seventh volume containing two issues with 18 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”[15] took place which was complemented with two seminars in 2016: “Reproducibility of Data-Oriented Experiments in e-Science”[16] and “Rethinking Experimental Methods in Computing”[17].

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.

[13] https://www.dagstuhl.de/dagman
[14] https://www.dagstuhl.de/darts
[15] https://www.dagstuhl.de/15452
[16] https://www.dagstuhl.de/16041
[17] https://www.dagstuhl.de/16111

Web Futures: Inclusive, Intelligent, Sustainable The 2020 Manifesto for Web Science (Dagstuhl Perspectives Workshop 18262)
Dagstuhl Manifestos, Volume 9, Issue 1, pp. 1–42, https://doi.org/10.4230/DagMan.9.1.1
based on Dagstuhl Perspectives Workshop 18262 https://www.dagstuhl.de/18262

Diversity in News Recommendation (Dagstuhl Perspectives Workshop 19482)
Dagstuhl Manifestos, Volume 9, Issue 1, pp. 43–61, https://doi.org/10.4230/DagMan.9.1.43
based on Dagstuhl Perspectives Workshop 19482 https://www.dagstuhl.de/19482

Fig. 4.1
Manifestos published in the 2021 volume of the journal *Dagstuhl Manifestos*.

Dagstuhl Publishing

OASIcs: OpenAccess Series in Informatics

Die *OASIcs*-Reihe[18] 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 2021 wurden 11 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.3.

OASIcs: OpenAccess Series in Informatics

The *OASIcs* series[18] 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 2021, Dagstuhl published 11 *OASIcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 4.3.

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

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

Fig. 4.2
OASIcs Editorial Board.

Vol. 82 | 2nd International Conference on Blockchain Economics, Security and Protocols (Tokenomics 2020) http://www.dagstuhl.de/dagpub/978-3-95977-157-3
Vol. 87 | Second Workshop on Next Generation Real-Time Embedded Systems (NG-RES 2021) http://www.dagstuhl.de/dagpub/978-3-95977-178-8
Vol. 88 | 12th Workshop on Parallel Programming and Run-Time Management Techniques for Many-core Architectures and 10th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms (PARMA-DITAM 2021) http://www.dagstuhl.de/dagpub/978-3-95977-181-8
Vol. 89 | 2nd International Conference of the DFG International Research Training Group 2057 – Physical Modeling for Virtual Manufacturing (iPMVM 2020) http://www.dagstuhl.de/dagpub/978-3-95977-183-2
Vol. 90 | 27th IFIP WG 1.5 International Workshop on Cellular Automata and Discrete Complex Systems (AUTOMATA 2021) http://www.dagstuhl.de/dagpub/978-3-95977-189-4
Vol. 91 | Second International Computer Programming Education Conference (ICPEC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-194-8
Vol. 92 | 4th International Symposium on Foundations and Applications of Blockchain 2021 (FAB 2021) http://www.dagstuhl.de/dagpub/978-3-95977-196-2
Vol. 93 | 3rd Conference on Language, Data and Knowledge (LDK 2021) http://www.dagstuhl.de/dagpub/978-3-95977-199-3
Vol. 94 | 10th Symposium on Languages, Applications and Technologies (SLATE 2021) http://www.dagstuhl.de/dagpub/978-3-95977-202-0
Vol. 95 | 3rd International Workshop on Formal Methods for Blockchains (FMBC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-209-9
Vol. 96 | 21st Symposium on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems (ATMOS 2021) http://www.dagstuhl.de/dagpub/978-3-95977-213-6

Fig. 4.3
OASIcs volumes published in 2021.

Dagstuhl Publishing

LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe[19] 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 Luca Aceto, Anca Muscholl, Dieter van Melkebeek, Catuscia Palamidessi und Thomas Schwentick sind 2021 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.

Luca Aceto und Anca Muscholl wurden in einem anonymen Wahlverfahren innerhalb des Herausgebergremiums für weitere 4 Jahre in das Gremium gewählt. Zudem wurden Faith Ellen, Daniel Král', Phillip Rogaway und Eva Rotenberg neu in das Gremium gewählt. Des weiteren wurde Luca Aceto im Oktober 2021 als Vorsitzender des Herausgebergremiums für zwei weitere Jahre bestätigt. Siehe auch Fig. 4.4.

In 2021 wurden Tagungsbände von 32 Konferenzen mit insgesamt 1333 Artikeln veröffentlicht; siehe Fig. 4.5 und 4.6.

Auch im zurückliegenden Jahr 2021 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 2021 bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischerweise 5 Jahre) eingegangen wurde. Zwei 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[19] 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 Luca Aceto, Anca Muscholl, Dieter van Melkebeek, Catuscia Palamidessi, and Thomas Schwentick ended in 2021. 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.

Luca Aceto and Anca Muscholl were voted in anonymous voting within the editorial board for another 4 years membership. In addition, Faith Ellen, Daniel Král', Phillip Rogaway and Eva Rotenberg were newly elected to the board. Furthermore, Luca Aceto was confirmed as chair of the editorial board for another two years in October 2021. See also Fig. 4.4.

The series published the proceedings of 32 major conferences with 1333 articles in total in 2021; see Fig. 4.5 and 4.6.

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2021 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 2021 for a cooperation covering several years (typically 5 years). Two of these conferences have submitted a proposal to LIPIcs for the first time. The other conferences have already cooperated with LIPIcs in the past.

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

Prof. Dr. Luca Aceto Reykjavik University, IS and GSSI, IT | tenure started in 2021 | Chair	Prof. Dr. Anca Muscholl University of Bordeaux, FR | tenure started in 2021
Prof. Dr. Christel Baier TU Dresden, DE	Prof. Dr. Luke Ong University of Oxford, GB
Prof. Dr. Mikolaj Bojanczyk University of Warsaw, PL	Dr. Catuscia Palamidessi INRIA, FR | tenure ended in 2021
Prof. Dr. Roberto Di Cosmo Inria and Université de Paris, FR	Prof. Dr. Phillip Rogaway University of California – Davis, US | tenure started in 2021
Prof. Dr. Faith Ellen University of Toronto, CA | tenure started in 2021	Prof. Dr. Eva Rotenberg Technical University of Denmark – Lyngby, DK | tenure started in 2021
Prof. Dr. Javier Esparza TU München, DE	Prof. Dr. Thomas Schwentick TU Dortmund, DE | tenure ended in 2021
Prof. Dr. Daniel Král' Masaryk University – Brno, CZ | tenure started in 2021	Prof. Raimund Seidel, Ph. D. Universität des Saarlandes – Saarbrücken, DE
Prof. Dr. Meena Mahajan Institute of Mathematical Sciences – Chennai, IN	Prof. Dieter van Melkebeek, Ph. D. University of Wisconsin-Madison, US | tenure ended in 2021

Fig. 4.4
LIPIcs Editorial Board.

Vol. 183 | 29th EACSL Annual Conference on Computer Science Logic (CSL 2021) http://www.dagstuhl.de/dagpub/978-3-95977-175-7
Vol. 184 | 24th International Conference on Principles of Distributed Systems (OPODIS 2020) http://www.dagstuhl.de/dagpub/978-3-95977-176-4
Vol. 185 | 12th Innovations in Theoretical Computer Science Conference (ITCS 2021) http://www.dagstuhl.de/dagpub/978-3-95977-177-1
Vol. 186 | 24th International Conference on Database Theory (ICDT 2021) http://www.dagstuhl.de/dagpub/978-3-95977-179-5
Vol. 187 | 38th International Symposium on Theoretical Aspects of Computer Science (STACS 2021) http://www.dagstuhl.de/dagpub/978-3-95977-180-1
Vol. 188 | 26th International Conference on Types for Proofs and Programs (TYPES 2020) http://www.dagstuhl.de/dagpub/978-3-95977-182-5
Vol. 189 | 37th International Symposium on Computational Geometry (SoCG 2021) http://www.dagstuhl.de/dagpub/978-3-95977-184-9
Vol. 190 | 19th International Symposium on Experimental Algorithms (SEA 2021) http://www.dagstuhl.de/dagpub/978-3-95977-185-6
Vol. 191 | 32nd Annual Symposium on Combinatorial Pattern Matching (CPM 2021) http://www.dagstuhl.de/dagpub/978-3-95977-186-3
Vol. 192 | 2nd Symposium on Foundations of Responsible Computing (FORC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-187-0
Vol. 193 | 12th International Conference on Interactive Theorem Proving (ITP 2021) http://www.dagstuhl.de/dagpub/978-3-95977-188-7
Vol. 194 | 35th European Conference on Object-Oriented Programming (ECOOP 2021) http://www.dagstuhl.de/dagpub/978-3-95977-190-0
Vol. 195 | 6th International Conference on Formal Structures for Computation and Deduction (FSCD 2021) http://www.dagstuhl.de/dagpub/978-3-95977-191-7
Vol. 196 | 33rd Euromicro Conference on Real-Time Systems (ECRTS 2021) http://www.dagstuhl.de/dagpub/978-3-95977-192-4
Vol. 197 | 16th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-198-6
Vol. 198 | 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021) http://www.dagstuhl.de/dagpub/978-3-95977-195-5
Vol. 199 | 2nd Conference on Information-Theoretic Cryptography (ITC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-197-9
Vol. 200 | 36th Computational Complexity Conference (CCC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-193-1
Vol. 201 | 21st International Workshop on Algorithms in Bioinformatics (WABI 2021) http://www.dagstuhl.de/dagpub/978-3-95977-200-6
Vol. 202 | 46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021) http://www.dagstuhl.de/dagpub/978-3-95977-201-3

Fig. 4.5
LIPIcs volumes published in 2021 – Part 1.

Dagstuhl Publishing

Vol. 203 | 32nd International Conference on Concurrency Theory (CONCUR 2021) http://www.dagstuhl.de/dagpub/978-3-95977-203-7
Vol. 204 | 29th Annual European Symposium on Algorithms (ESA 2021) http://www.dagstuhl.de/dagpub/978-3-95977-204-4
Vol. 205 | 27th International Conference on DNA Computing and Molecular Programming (DNA 27) http://www.dagstuhl.de/dagpub/978-3-95977-205-1
Vol. 206 | 28th International Symposium on Temporal Representation and Reasoning (TIME 2021) http://www.dagstuhl.de/dagpub/978-3-95977-206-8
Vol. 207 | Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021) http://www.dagstuhl.de/dagpub/978-3-95977-207-5
Vol. 208 | 11th International Conference on Geographic Information Science (GIScience 2021) – Part II http://www.dagstuhl.de/dagpub/978-3-95977-208-2
Vol. 209 | 35th International Symposium on Distributed Computing (DISC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-210-5
Vol. 210 | 27th International Conference on Principles and Practice of Constraint Programming (CP 2021) http://www.dagstuhl.de/dagpub/978-3-95977-211-2
Vol. 211 | 9th Conference on Algebra and Coalgebra in Computer Science (CALCO 2021) http://www.dagstuhl.de/dagpub/978-3-95977-212-9
Vol. 212 | 32nd International Symposium on Algorithms and Computation (ISAAC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-214-3
Vol. 213 | 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2021) http://www.dagstuhl.de/dagpub/978-3-95977-215-0
Vol. 214 | 16th International Symposium on Parameterized and Exact Computation (IPEC 2021) http://www.dagstuhl.de/dagpub/978-3-95977-216-7

Fig. 4.6
LIPIcs volumes published in 2021 – Part 2.

ECRTS | Euromicro Conference on Real-Time Systems accepted for 2022-2026 (Re-evaluation)
ESA | European Symposium on Algorithms accepted for 2021-2025 (Re-evaluation)
FUN | International Conference on Fun with Algorithms accepted for 2022-2026 (Re-evaluation)
ITCS | Innovations in Theoretical Computer Science Conference accepted for 2022-2026 (Re-evaluation)
MFCS | Mathematical Foundations of Computer Science accepted for 2021-2025 (Re-evaluation)
SAND | Symposium on Algorithmic Foundations of Dynamic Networks accepted for 2022-2024
SAT | International Conference on Theory and Applications of Satisfiability Testing accepted for 2022-2026
SEA | International Symposium on Experimental Algorithms accepted for 2022 (Re-evaluation)

Fig. 4.7
Conferences accepted in 2021 for publication in LIPIcs.

Dagstuhl Publishing

LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*[20] veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. Die Zeitschrift wurde 2012 gegründet und nahm 2013 ihren Betrieb auf. 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)*[21] der Fachgesellschaft *European Design and Automation Association (EDAA)*[22] 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 2021 wurde eine Ausgabe von *LITES* mit drei Artikeln veröffentlicht.

LITES: Leibniz Transactions on Embedded Systems

The *LITES*[20] 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)*[21] of the *European Design and Automation Association (EDAA)*[22]. 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 2021, one issue of *LITES* containing three articles was published.

[20] https://www.dagstuhl.de/lites
[21] http://www.emsig.net/
[22] https://www.edaa.com/

Prof. Alan Burns, DPhil University of York, UK | Editor-in-Chief	Prof. Bashir Al Hashimi University of Southampton, UK
Prof. Sang Lyul Min, Ph. D. Seoul National University, KR	Prof. Dr. Martin Fränzle Carl von Ossietzky University Oldenburg, DE
Prof. Dr. Marco di Natale Scuola Superiore Santa Anna, IT	Prof. Dr. Samarjit Chakraborty Technical University Munich, DE
Dr. Virginie Wiels ONERA, FR	Prof. Dr. Gernot Heiser University of New South Wales, AT
Prof. Karl-Erik Arzén, Ph. D. Lund University, SE	Prof. Dr. Lothar Thiele ETH Zürich, CH
Prof. Steve Goddard, Ph. D. University of Nebraska-Lincoln, US	Dr. Neil Audsley University of York, UK
Prof. Dr. Axel Jantsch Technical University of Vienna, AT	Prof. Sanjoy Baruah, Ph. D. University of North Carolina at Chapel Hill, US

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[23] 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, was eine gute Sichtbarkeit der Publikationen garantiert.

LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*[24] 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[23] 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*[24] 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.

[23] https://www.scopus.com
[24] 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 im Arbeitskreis Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Single Source Publishing“[25] mit organisiert, welcher bereits der sechste Workshop in Folge seit 2013 ist. Der Workshop fand am 14. April 2021 in virtueller Form statt. Dr. Michael Wagner, Mitglied des wissenschaftlichen Stabs von Schloss Dagstuhl, ist seit Oktober 2020 stellvertretender Sprecher des Arbeitskreises.

Publikationsserver: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)[26] werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative[27] berücksichtigt, insbesondere werden alle nötigen Metadaten zu jeder Publikation gespeichert und die Langzeitverfügbarkeit sichergestellt. Die Online-Publikationen sind zitierfähig und stehen einer großen Leserschaft zur Verfügung.

Einreichungssystem: DSub

Im Frühjahr 2019 wurde ein von Dagstuhl entwickeltes Einreichungssystem namens DSub in Betrieb genommen. Mit diesem System werden seitdem alle Einreichungen zu den Reihen LIPIcs und OASIcs entgegengenommen. Unter anderem wurde mit diesem 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 (DNB)[28] 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.[29]
- 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.[30]

Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Single Source Publishing”[25] was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop took place as a virtual event on April 14, 2021. Dr. Michael Wagner, a member of the scientific staff of Schloss Dagstuhl, is the elected deputy spokesperson of the working group since October 2020.

Publication Server: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)[26]. The general guidelines of the Dublin Core initiative[27] 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 (DNB)[28] 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[29].
- 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.[30]

[25] https://www.dagstuhl.de/fileadmin/dagpub/journalmanagement-leibniz/2021-04-workshop/
[26] https://www.dagstuhl.de/drops
[27] http://dublincore.org/
[28] https://www.dnb.de/DE/Professionell/Erhalten/erhalten_node.html#sprg209698
[29] https://subs.emis.de/LIPIcs/
[30] http://vesta.informatik.rwth-aachen.de/Dagstuhl/

5 Resonanz
Feedback

Resonanz zu Seminaren und Workshops

5.1

Feedback on Seminars and Workshops

Resonanz von Teilnehmern

Schloss Dagstuhl bekommt viel Lob von seinen Gästen, meistens in mündlicher Form, wenn die Gäste nach einer intensiven Seminarwoche das Schloss verlassen. Manche Gäste nehmen sich jedoch auch die Zeit, uns nachträglich zu schreiben und ihre Eindrücke mit uns zu teilen.

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.

Email from Simon Fowler (University of Glasgow, GB)

21372 – Behavioural Types: Bridging Theory and Practice | Dagstuhl Seminar | https://www.dagstuhl.de/21372

I attended Seminar 21372 last week and just wanted to write to express my distinct gratitude for my time at Schloss Dagstuhl. The staff were incredibly friendly and helpful, my room was spacious and immaculate when I arrived, and the food was constantly delicious. Thank you so much – attending my first Dagstuhl seminar was an experience I will not forget.

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 wir 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.

Post card from research group meeting 21103

21103 – Explainable Artificial Intelligence | Meeting | https://www.dagstuhl.de/21103

Die Teilnehmer der Projekt Treffens *Anwendungen formaler Wissenschaften: Explainable Artificial Intelligence* waren beeindruckt und berührt von der außerordentlichen Freundlichkeit der gesamten Belegschaft des Leibniz-Zentrums für Informatik. Für die umfassende Unterstützung und die exzellente Atmosphäre bedanken wir und herzlich!

The organizers of Dagstuhl Seminar 21293, Dagstuhl Report

21121 – Computational Complexity of Discrete Problems | Dagstuhl Seminar | https://www.dagstuhl.de/21121

We would also like to especially thank the Dagstuhl staff for their cooperation in the current challenging circumstances, their encouragement for going ahead with an online event, and their unstinted help with organizational matters.

The organizers of Dagstuhl Seminar 21293, Dagstuhl Report

21293 – Parameterized Complexity in Graph Drawing | Dagstuhl Seminar | https://www.dagstuhl.de/21293

Schloss Dagstuhl was the perfect place for hosting a seminar like this. The unique scientific atmosphere and the historic building provided not only all the room we needed for our program and the working groups, but also plenty of opportunities for continued discussions and socializing outside the official program, especially in these difficult times during the COVID-19 pandemic with all participants being eager to meet and do research together in real life. On behalf of all participants, the organizers want to express their deep gratitude to the entire Dagstuhl staff for their outstanding support and service accompanying this seminar.

Email from the organizers of Dagstuhl Seminar 21471

21471 – Geometric Modeling: Interoperability and New Challenges | Dagstuhl Seminar | https://www.dagstuhl.de/21471

[. . .]Every thing is so easy and well organised in Dagstuhl, that we are really dedicated to the scientific happening, [. . .]

Resonanz in Sozialen Netzwerken

Mehr und mehr Gäste nutzen die Möglichkeiten des Webs wie Twitter und Blogs um 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.

Dirk Fahland (TU Eindhoven, NL)
Twitter | https://twitter.com/dfahland/status/1452890447362338817

Seriously, let's build a network of Dagstuhl's all over the planet. That's why we are going to conferences. Spend one week of in depth discussion.

Bernardo Toninho (NOVA School of Science and Technology – Lisbon, PT)
Twitter | https://twitter.com/bernpton/status/1368016160059441155

Maybe this is what the whole vocation aspect of academia means, go and live in a remote CS german castle. That would do…

Eren Yildiz (University of Umeå, SE)
Google maps | https://goo.gl/maps/hcSsJ8qhjL5T1R3G7

Amazing place. Great food. I can't think a better place to do research.

Resonanz im Fragebogen

Jeder Teilnehmer erhält von uns einen Fragebogen zur Evaluation des vom Teilnehmer besuchten Dagstuhl-Seminars oder Dagstuhl-Perspektiven-Workshops. Durch diese anonyme 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.

21152 – Multi-Level Graph Representation for Big Data Arising in Science Mapping | Dagstuhl Seminar | https://www.dagstuhl.de/21152

Thanks for making this possible in spite of corona! The rapid testing facilities are great!

21201 – Serverless Computing | Dagstuhl Seminar | https://www.dagstuhl.de/21201

The Dagstuhl Seminar is unique and I am always returning to it with great pleasure. Please keep offering this level of support

21271 – Computational Proteomics | Dagstuhl Seminar | https://www.dagstuhl.de/21271

Keep it going. I enjoy Dagstuhl seminars a lot both scientifically and socially.

21293 – Parameterized Complexity in Graph Drawing | Dagstuhl Seminar | https://www.dagstuhl.de/21293

Thanks to the seminar organizers and the Dagstuhl staff – participating in such seminars is the usually the best week of my year!

21292 – Scalable Handling of Effects | Dagstuhl Seminar | https://www.dagstuhl.de/21292

This was the best virtual event that I attended in the past 1.5 years. Big thanks to the organizers!

21293 – Parameterized Complexity in Graph Drawing | Dagstuhl Seminar | https://www.dagstuhl.de/21293

Thanks everyone for making a safe on-site Dagstuhl seminar possible in pandemic times!

21351 – Universals of Linguistic Idiosyncrasy in Multilingual Computational Linguistics | Dagstuhl Seminar | https://www.dagstuhl.de/21351

thanks a lot, it was a wonderful experience! scientific and networking wise

21442 – Ensuring the Reliability and Robustness of Database Management Systems | Dagstuhl Seminar | https://www.dagstuhl.de/21442

I would like to thank Dagstuhl staff for enabling the in-person seminars during this difficult time

21441 – Adaptive Resource Management for HPC Systems | Dagstuhl Seminar | https://www.dagstuhl.de/21441

Thank you (to all the staff) for continuing this remarkable place in these difficult times.

21451 – Managing Industrial Control Systems Security Risks for Cyber Insurance | Dagstuhl Seminar | https://www.dagstuhl.de/21451

Thanks for the extensive feedback possibility!

21462 – Foundations of Persistent Programming | Dagstuhl Seminar | https://www.dagstuhl.de/21462

I'm honored to participate in Dagstuhl. I met legendary figures, though remotely, which is a unique experience for me as a junior researcher. I wish maybe next time we should encourage physical attendance more.

21492 – Representing and Solving Spatial Problems | Dagstuhl Seminar | https://www.dagstuhl.de/21492

I really value that Schloss Dagstuhl is family friendly. Thanks for everything

21293 – Parameterized Complexity in Graph Drawing | Dagstuhl Seminar | https://www.dagstuhl.de/21293

Food in Dagstuhl has always been very good, but Corona made it even better. The only think that I'm missing is the traditional Dagstuhl muesli ;-)

21381 – Conversational Agent as Trustworthy Autonomous System (Trust-CA) | Dagstuhl Seminar | https://www.dagstuhl.de/21381

I would like to thank the kitchen staff for their work in preparing and presenting the meals in such a lovely way.

21381 – Conversational Agent as Trustworthy Autonomous System (Trust-CA) | Dagstuhl Seminar | https://www.dagstuhl.de/21381

Dagstuhl staff: they are super friendly and helpful. Quality of meals: excellent quality. Kitchen staff were very kind. All worked very nicely!

21452 – Unambiguity in Automata Theory | Dagstuhl Seminar | https://www.dagstuhl.de/21452

It is amazing how much care is devoted to offer vegetarians convenient meals! Many thanks.

21071 – Scalable Data Structures | Dagstuhl Seminar | https://www.dagstuhl.de/21071

That this Dagstuhl, despite of the pandemic, was not canceled but organized as an online Dagstuhl. This worked out quite well and was a very good decision by the organisers but on the other hand showed how much we should appreciate the possibilities Dagstuhl offers under regular circumstances.

21152 – Multi-Level Graph Representation for Big Data Arising in Science Mapping | Dagstuhl Seminar | https://www.dagstuhl.de/21152

best: that a seminar happened at all after such a long time, and I was pleasantly surprised that a hybrid format was useful; the evening events were a great idea! worst: that it was impossible to attend in person due to quarantine regulations; as expected remote participation was more in conflict with other duties at the home institution (but many participants still managed to attend as much as possible).

21192 – Approaches and Applications of Inductive Programming | Dagstuhl Seminar | https://www.dagstuhl.de/21192

unfortunately, it was only online; but nevertheless, discussions after talks as well as between participants (at gather.town) were lively, focussed, and inspiring

21243 – Compute-First Networking | Dagstuhl Seminar | https://www.dagstuhl.de/21243

The best part was being part of fascinating discussions with many brilliant people. The worst aspect was the inability to on-site due to the Covid-19 restrictions.

21262 – Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users | Dagstuhl Seminar | https://www.dagstuhl.de/21262

Best: interacting and exchanging ideas with people from academia and industry

21261 – Quantum Complexity: Theory and Application | Dagstuhl Seminar | https://www.dagstuhl.de/21261

The best thing was to finally be able to meet people in person again and to talk over a glass of wine in the evening and enjoy music together. Such informal discussions and interactions is something we all have missed very much over the last year and a half.

21271 – Computational Proteomics | Dagstuhl Seminar | https://www.dagstuhl.de/21271

Best: The opportunity to talk with, listen to and workshop with colleagues on an extended topic. It provided an excellent opportunity to get different perspectives – far more than you'd easily be able to achieve at a normal conference Worst: I unfortunately was unable to attend in person because of the pandemic. The hybrid format worked well (better than I expected), but you do feel you miss out on the face to face contact. Hopefully next time I'll be able to attend in person!

21293 – Parameterized Complexity in Graph Drawing | Dagstuhl Seminar | https://www.dagstuhl.de/21293

The best aspect is the friendly and interdisciplinary atmosphere. I cannot think of negative remarks.

21342 – Identifying Key Enablers in Edge Intelligence | Dagstuhl Seminar | https://www.dagstuhl.de/21342

Best aspects: The seminar helped me understand the status of the field and future possibilities in terms of research directions in industry and academia. The seminar helped me connect with the cutting-edge researchers worldwide. Worst aspects: Virtual attendance made it hard to have peer to peer social interactions.

21351 – Universals of Linguistic Idiosyncrasy in Multilingual Computational Linguistics | Dagstuhl Seminar | https://www.dagstuhl.de/21351

The opportunity to meet people from across the research community in a relaxed environment without the added stresses that a conference can add (presentations, attending talks etc). The only drawback was that it was online, but that's an uncontrollable factor this year!

21362 – Structure and Learning | Dagstuhl Seminar | https://www.dagstuhl.de/21362

While the hybrid format is better than a fully virtual meeting (at least for the people on-site), it still is far from a fully physical meeting.

21352 – Higher-Order Graph Models: From Theoretical Foundations to Machine Learning | Dagstuhl Seminar | https://www.dagstuhl.de/21352

Best: as a PhD student I met some important people. Worst: it's too short.

21361 – Extending the Synergies Between SAT and Description Logics | Dagstuhl Seminar | https://www.dagstuhl.de/21361

I was, in the end, unable to attend in person, due to the pandemic. This experience has just reinforced to me how important and valuable in-person Dagstuhl seminars are.

21371 – Integrated Deduction | Dagstuhl Seminar | https://www.dagstuhl.de/21371

Participants are the world experts, there's no better crowd to discuss state of the art. Schloss Dagstuhl provides an extraordinary environment to concentrate, search and understand science.

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. Zudem wurde 2021 eigens eine Nutzerbefragung zu dblp durchgeführt.

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. Further comments have been collected in a dedicated dblp user survey in 2021.

Resonanz in Sozialen Netzwerken

Social Media Feedback

Josip Matak, Memgraph Ltd.

Blog Article | https://memgraph.com/blog/influencers-among-computer-scientists

There is no better source to determine who the most influential computer scientists are than looking into DBLP - computer science bibliography. It is a website that collects and hosts the largest amount of articles published by various computer scientists in different fields of expertise. There you can find an awesome dataset that contains information about paper citations.

Pasquale Minervini, University College London, UK

Twitter | https://twitter.com/PMinervini/status/1355061964305821696

I just link to the @dblp_org bibtex on my website so everyone is happy

Gergely Neu, Pompeu Fabra University, Spain

Twitter | https://twitter.com/neu_rips/status/1367064171565748224

DBLP >>>> Gscholar at least in terms of quality of bibtex files (e.g., more consistent conference names, page numbers..)

Peter Suber, Harvard Library, MA, USA

Twitter | https://twitter.com/petersuber/status/1369649187168661507

Nice call out to the "incredible open research ecosystem built…by #OpenAccess advocates" — citing DOAJ, Biodiversity Heritage Library, Unpaywall, CORE, CiteseerX, MS Academic, Semantic Scholar, ROAD, SHERPA/ROMEO, JURN, Wikidata, Crossref, Datacite, J-STAGE, Pubmed, & dblp.

Jana Dunfield, Queen's University, Kingston, ON, Canada

Twitter | https://twitter.com/etrolleybus/status/1377478998796734464

#TransNameChangePolicy status:
DBLP: emailed 2020-12-10, all fixed next day
CMU CS: fixed non-thesis TRs, won't fix thesis without legal name change
ACM: emailed their form 2021-02-22, no progress
arXiv: requested 2021-03-23, completed 03-25, won't fix self-citations
Elsevier: emailed 2021-03-29, replied 03-30, "will take $\leq$2 weeks"
CUP: emailed 2021-03-29, no reply
EPTCS: not bothering with this yet, the arXiv copies get ranked higher than their secret copies
Springer: …yeah, not going there

Georg Weissenbacher, TU Vienna, Austria

Twitter | https://twitter.com/gweissenbacher/status/1377557015808053250

I very much prefer DBLP - I rarely ever have to worry about it, whereas ORCID requires constant grooming and still looks ugly. I never felt the urge to look up somebodies ORCID record.

Tianyin Xu, University of Illinois Urbana-Champaign, IL, USA
Twitter | https://twitter.com/tianyin_xu/status/1403030045279596544

How can @dblp_org be so incredibly responsive and willing to help?!!! Thank you!!!

Adam Elmachtoub, Columbia University, NY, USA
Twitter | https://twitter.com/Adam235711/status/1378520179009847296

I love how @dblp_org displays my work. I wish I could link to my dblp page from my website, with the "Informal Publications" box already unchecked when you land. Anyway to do that?

Nadia Polikarpova, UC San Diego, CA, USA
Twitter | https://twitter.com/polikarn/status/1383874158287736832

"Have you googled your own name in the past year?"
Does "<my name> dblp" count? That's how I get bibtexs for my own papers :)

Robiul Islam, Innopolis University, Russia
Twitter | https://twitter.com/connect2robiul/status/1386783865126244357

I am extremely happy that my publication cited on @dblp_org very first time.

Johann Dreo, Institut Pasteur, France
Twitter | https://twitter.com/j_dreo/status/1389920885914619908

What to do when the JabRef extension cannot find/parse the structured metadata about the reference? The best option is usually to try on DBLP, which have both a comprehensive database (in Computer Science) AND a good BiBTeX export.

Sam Power, University of Bristol, UK
Twitter | https://twitter.com/sp_monte_carlo/status/1377663611057889280

after reading this tweet and some of the replies about a month ago, i was convinced to start using DBLP (for ML conference papers, at least) instead of google scholar, and i have been very satisfied with the results.

Adam Elmachtoub, Columbia University, NY, USA
Twitter | https://twitter.com/Adam235711/status/1408468107862151182

Small victory.. I reached out to the editorial team @informs journal on applied analytics (IJAA) to let them know @dblp_org was still using the old name (interfaces). Fixed in just a week's time!

Dominik Bork, TU Vienna, Austria
Twitter | https://twitter.com/BorkDominik/status/1435337596113235969

Just saw, for the first time, this new nice visualization of research output integrated with @dblp_org It color-codes publication of each year grouped by publication type.

Michele Lanza, USI, Switzerland
Twitter | https://twitter.com/lanzamichele/status/1438010595471761409

Small idea with some funny consequences: Have google scholar sync with DBLP and only keep the publications on g scholar that are actually listed on DBLP

Bastian Tenbergen, SUNY at Oswego, NY, USA
Twitter | https://twitter.com/BTenbergen/status/1389547911797739522

Dating ppl be like "show me their Twitter // show me their Insta" and I'm like, darling... Show me their @ResearchGate // show me their @dblp_org

Hugo Lima, Federal University of Pará, Brazil
Twitter | https://twitter.com/limahugob/status/1389003102422962176

Love you, @dblp_org <3

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.

anonymous survey comment

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

I think Dagstuhl Publishing is doing excellent job regarding guidelines and instructions. The LaTeX style is very elegant.

anonymous survey comment

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

I thought it all went very smoothly and the style is very elegant. Congratulations, and thank you very much for this great service!

anonymous survey comment

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

I liked the way you encourage authors to publish full versions of their paper (and provide access to code). This is very beneficial for the scientific process. Overall, a very good experience. Quite smooth, and with lot of attention to quality.

anonymous survey comment

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

The new DSUB is *amazing*! I remember doing it for 2016 and the new submission system is far better. I mean, I am *really* impressed at how good it is.

Resonanz in Sozialen Netzwerken

Social Media Feedback

Jesper Agdakx

Twitter | https://twitter.com/agdakx/status/1283329505331945472

Huge props to @dagstuhl for not only providing open-access journals with reasonable author fees but also building the first paper submission site that's actually a pleasure to use. @TheOfficialACM, @CambridgeUP & others could really learn something.

6 Die Seminare in 2021
The 2021 Seminars

Applications, Interdisciplinary Work

- Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments (21232)
- Computational Proteomics (21271)
- Towards Climate-Friendly Internet Research (21272)
- Digital Disinformation: Taxonomy, Impact, Mitigation, and Regulation (21402)
- Machine Learning in Sports (21411)
- Managing Industrial Control Systems Security Risks for Cyber Insurance (21451)

Artificial Intelligence, Computational Linguistics

- Autonomous Agents on the Web (21072)
- Transparency by Design (21231)
- Universals of Linguistic Idiosyncrasy in Multilingual Computational Linguistics (21351)
- Structure and Learning (21362)
- Conversational Agent as Trustworthy Autonomous System (Trust-CA) (21381)
- Representing and Solving Spatial Problems (21492)

Cryptography, Security, Privacy

- Quantum Cryptanalysis (21421)
- Rigorous Methods for Smart Contracts (21431)

Databases, Information Retrieval, Machine Learning, Data Mining

- Probabilistic Numerical Methods – From Theory to Implementation (21432)
- Ensuring the Reliability and Robustness of Database Management Systems (21442)

Data Structures, Algorithms, Complexity

- Scalable Data Structures (21071)
- Computational Complexity of Discrete Problems (21121)
- Temporal Graphs: Structure, Algorithms, Applications (21171)
- Computational Geometry (21181)
- Quantum Complexity: Theory and Application (21261)
- Data Structures for Modern Memory and Storage Hierarchies (21283)
- Parameterized Complexity in Graph Drawing (21293)
- Matching Under Preferences: Theory and Practice (21301)
- Coalition Formation Games (21331)
- Higher-Order Graph Models: From Theoretical Foundations to Machine Learning (21352)
- Sparsity in Algorithms, Combinatorics, and Logic (21391)
- Descriptive Set Theory and Computable Topology (21461)

Distributed Computation, Networks, Architecture, Systems

- Serverless Computing (21201)
- Compute-First Networking (21243)
- Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users (21262)
- Understanding I/O Behavior in Scientific and Data-Intensive Computing (21332)
- Identifying Key Enablers in Edge Intelligence (21342)
- Adaptive Resource Management for HPC Systems (21441)
- Foundations of Persistent Programming (21462)

Geometry, Image Processing, Graphics, Visualization

- Multi-Level Graph Representation for Big Data Arising in Science Mapping (21152)
- Visualization of Biological Data – From Analysis to Communication (21401)
- Geometric Modeling: Interoperability and New Challenges (21471)

Software Technology, Programming Languages

- Approaches and Applications of Inductive Programming (21192)
- Scalable Handling of Effects (21292)
- Approximate Systems (21302)
- Behavioural Types: Bridging Theory and Practice (21372)
- Secure Compilation (21481)

Verification, Logic, Formal Methods, Semantics

- Extending the Synergies Between SAT and Description Logics (21361)
- Integrated Deduction (21371)
- Unambiguity in Automata Theory (21452)
- Geometric Logic, Constructivisation, and Automated Theorem Proving (21472)

6.1 Scalable Data Structures

Organizers: Gerth Stølting Brodal, John Iacono, Markus E. Nebel, and Vijaya Ramachandran

Seminar No. 21071

Date: February 14–19, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.1.1

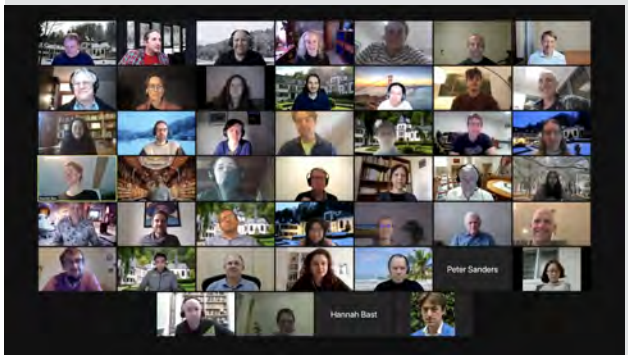

Remote Participants: Kunal Agrawal, Elena Arseneva, Hannah Bast, Ioana Oriana Bercea, Guy E. Blelloch, Gerth Stølting Brodal, Rezaul Chowdhury, Alexander Conway, Martin Dietzfelbinger, Anne Driemel, Amalia Duch Brown, Faith Ellen, Guy Even, Rolf Fagerberg, Martin Farach-Colton, Afton Noelle Geil, Inge Li Gørtz, Monika Henzinger, John Iacono, Riko Jacob, Rob Johnson, Valerie King, Tsvi Kopelowitz, László Kozma, Moshe Lewenstein, Andrea Lincoln, Quanquan C. Liu, Ulrich Carsten Meyer, Ian Munro, Markus E. Nebel, Eunjin Oh, John Owens, Rotem Oshman, Seth Pettie, Vijaya Ramachandran, Rajeev Raman, Eva Rotenberg, Peter Sanders, Robert Sedgewick, Julian Shun, Francesco Silvestri, Nodari Sitchinava, Tatiana Starikovskaya, Jesper Steensgard, Yihan Sun, Robert Endre Tarjan, Manoharan Vignesh, Sebastian Wild

About the seminar

Scalable data structures form the backbone for computing: Computing is about processing, exchanging, and storing data. The organization of data profoundly influences the performance of accessing and manipulating data. By optimizing the way data is stored, performance can be improved by several orders of magnitude when data scales. This Dagstuhl Seminar brought together researchers from several research directions to illuminate solutions to the scalability challenge of data structures. The seminar was the 14th in a series of loosely related Dagstuhl Seminars on data structures. Due to the ongoing Covid-19 pandemic the seminar was purely virtual.

Topics

The presentations covered both advances in classic data structure fields, as well as insights addressing the scalability of computing for different models of computation.

Classic data structure questions on dictionaries, hashing, filters, and heaps were the topic of several talks. Wild (Section 4.8 of the full report) considered new pointer based search trees, Kozma (Section 4.24 of the full report) discussed algorithms related to self-adjusting trees and heaps, Bercea (Section 4.4 of the full report) presented results for dictionaries and filters, Even (Section 4.21 of the full report) discussed dynamic stable perfect hashing, and Farach-Colton (Section 4.2 of the full report) presented results for succinct stable hash tables. Johnson (Section 4.25 of the full report) presented the vector quotient filter based on Robin Hood hashing and designed to exploit SIMD instructions. An external memory dictionary was presented by Conway (Section 4.20 of the full report) who presented the SplinterDB key-value store for NVMe solid state drives.

For data structure problems on strings, Gørtz (Section 4.9 of the full report) discussed the support for random access in compact representations of strings, and Starikovskaya (Section 4.22 of the full report) dictionary look-ups with mismatches.

Data structures for storing and querying static and dynamic graphs were the topic of a sequence of talks. Pettie (Section 4.1 of the full report) considered trade-offs between space usage and query time for supporting exact distance queries in planar graphs. Rotenberg (Section 4.5 of the full report) considered planarity testing of dynamic graphs under the insertion and deletion of edges with polylogarithmic update time. Kopelowitz (Section 4.6 of the full report) considered maintaining the orientation of edges in dynamic forests under the insertion of edges, guaranteeing low out degree of all nodes. Henzinger (Section 4.16 of the full report) presented an algorithm for maintaining a $(\Delta + 1)$-vertex coloring of a graph with maximal degree Δ with constant time edge insertions and deletions. Bast (Section 4.27 of the full report) gave a demonstration of an implementation of algorithms for real-time searching knowledge graphs with billions of edges.

Parallel algorithms for problems on graphs were addressed in multiple talks. Liu (Section 4.10 of the full report) considered a parallel algorithm for counting triangles (cliques of size three) in graphs under batched updates of edge insertions and deletions, and Blelloch (Section 4.15 of the full report) considered parallel batched dynamic algorithms for the minimum spanning tree and minimum cut problems. Sun (Section 4.13 of the full report) considered a parallel algorithm for the single source shortest path problem using lazy batched priority queues. Shun (Section 4.3 of the full report) considered a parallel index-based algorithm for graph clustering and an approximation algorithm using locality-sensitive hashing.

Computational models supporting massive parallelism, like GPUs and TCUs, were addressed in talks by Owens (Section 4.11 of the full report) who considered open-addressing hashing on GPUs, by Geil (Section 4.18 of the full report) who consider how to solve the maximum clique problem on GPUs, and by Silvestri (Section 4.14 of the full report) who addressed similarity

search with tensor core units. Sitchinava & Jacob (Section 4.19 of the full report) in their joint talk, considered the power of the atomic and non-atomic versions of the parallel fork-join model. Sanders (Section 4.12 of the full report) considered how to execute MapReduce computations robustly and efficiently on realistic distributed-memory parallel machines.

Ellen (Section 4.17 of the full report) considered labelling schemes for networks supporting distributed deterministic radio broadcast using labels of constant-length at the nodes of the network. Lincoln (Section 4.23 of the full report) presented new techniques for proving fine-grained average-case hardness results, and Fagerberg (Section 4.26 of the full report) considered the fragile complexity of adaptive algorithms. Finally, Driemel (Section 4.7 of the full report) considered approximate-near-neighbor data structures for time series under the continuous Fréchet distance.

Final Thoughts

The organizers would like to thank the Dagstuhl team for their continuous support and allowing this seminar to happen as a purely virtual Dagstuhl Seminar. They also thank all participants for their contributions to this seminar.

Previous seminars in the series had few female participants. A focus for this seminar was to significantly increase the female attendance. 50 % of the invited participants were female, resulting in a 38 % female attendance.

Even though the seminar was challenged by the different time zones of the participants, on average 37 of the 48 participants participants attended the talks, and all talks were attended by at least 30 partcipants. In the post-seminar survey it was appreciated that the seminar took place as a virtual seminar instead of being cancelled, but it was also stated that the virtual format can never be as productive as an in-person seminar and showed how much we should appreciate the possibilities Dagstuhl offers under regular circumstances.

6.2 Autonomous Agents on the Web

Organizers: Olivier Boissier, Andrei Ciortea, Andreas Harth, and Alessandro Ricci
Seminar No. 21072

Date: February 14–19, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.1.24

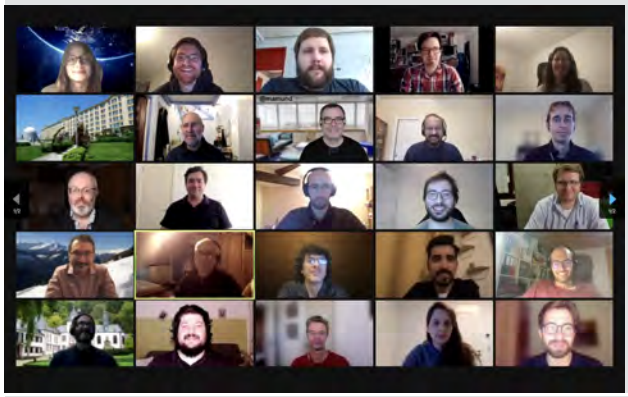

Remote Participants: Cleber Jorge Amaral, Mike Amundsen, Olivier Boissier, Rafael H. Bordini, Arne Bröring, Samuele Burattini, Jean-Paul Calbimonte, Victor Charpenay, Samuel Christie, Andrei Ciortea, Olivier Corby, Angelo Croatti, Jérôme Euzenat, Matthias Farnbauer-Schmidt, Fabien Gandon, Dominique Guinard, Andreas Harth, James A. Hendler, Jomi Fred Hübner, Tobias Käfer, Lalana Kagal, Timotheus Kampik, Sabrina Kirrane, Ege Korkan, Matthias Kovatsch, Maxime Lefrançois, Adnane Mansour, Simon Mayer, Mahda Noura, Julian Padget, Cesare Pautasso, Terry Payne, María Poveda-Villalón, Dave Raggett, Alessandro Ricci, Sebastian Schmid, Daniel Schraudner, Ronny Seiger, Munindar P. Singh, Ruben Taelman, Valentina Tamma, Danai Vachtsevanou, Ruben Verborgh, Erik Wilde, Antoine Zimmermann

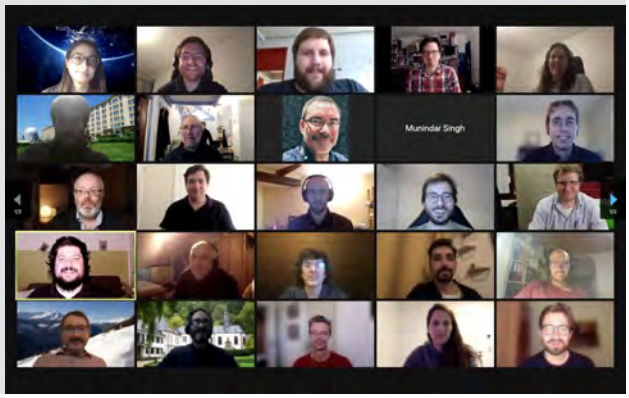

The vision of autonomous agents on the Web is almost as old as the Web itself: in his keynote at WWW'94[31], Sir Tim Berners-Lee was noting that documents on the Web describe real objects and relationships among them, and if the semantics of these objects are represented explicitly then machines can *browse through* and *manipulate* reality.[32] These ideas were published under the Semantic Web vision in 2001 [1]. Yet in 2007, after having spent the better half of a decade advancing this vision, James Hendler was looking back to conclude that most ideas in the original article were already seeing widespread deployment on the Web except for agent-based systems – and raised the question: "where are all the intelligent agents?" [2].

This question is yet to be addressed. On today's Web, we as *human agents* are often assisted by invisible software agents, such as crawlers used by search engines to navigate and index Web pages, agents that curate online content produced by people (e.g., Wikipedia's content agents), and recommender systems used all over the Web to generate more links and navigation paths (e.g., suggestions of related Web pages). In our everyday lives, we are assisted by more visible agents, such as Amazon's Alexa, Google Duplex, or Apple's Siri. Some of these agents may already use various AI methods (learning, reasoning, etc.), but they are specialized for narrow tasks and constrained to silos dictated by company ecosystems. We have yet to see more autonomous, cooperative, and long-lived agents on the Web [3] – the intelligent agents in James Hendler's question. We believe this decade-old question is now more relevant than ever before in the context of recent developments in three areas of research: (i) *Web Architecture and the Web of Things*, (ii) *Semantic Web and Linked Data*, and (iii) *Autonomous Agents and Multi-Agent Systems*.

The primary objective of this 5-day seminar was to create a network of senior and young researchers who can revisit and align the relevant research threads across the three targeted areas. The seminar was a blend of invited talks, live demonstrators,

[31] The First International Conference on the World-Wide Web, CERN, 25-27 May 1994.
[32] Sir Tim Berners-Lee, *The Future of the Web*, WWW'94: https://videos.cern.ch/record/2671957, accessed: 07.05.2021.

Fig. 6.1
The traditional end-of-seminar photo on the stairs of Virtual Schloss Dagstuhl.

and group work. Some of the overarching research questions discussed during the seminar included (not an exhaustive list): How to design software agents able to achieve their tasks through flexible autonomous use of hypermedia? How to design hypermedia-based environments that support autonomous behavior? How to design, represent, and reason about interactions among autonomous agents, people, and any other resources on the Web? How to design and govern communities of autonomous agents and people on the Web?

A number of follow-up steps were already taken to continue the discussions and to further consolidate the community. Most recently, several participants submitted a challenge proposal that was accepted at the *20th International Semantic Web Conference (ISWC 2021)*: the *All The Agents Challenge (ATAC)* will take place in October 2021.[33] In addition, a one-day follow-up event was scheduled for July 9, 2021, and several participants offered to contribute to a shared "live" demonstrator space that would allow to integrate technologies and to try out new ideas across the targeted research areas (see also the working group report in Section 6.1 of the full report).

The Virtual Seminar Format

The seminar was organized in a fully virtual format due to the COVID-19 pandemic and brought together 45 participants across 4 different time zones. The seminar's schedule was designed for focused online interaction and to preserve to some extent the social dimension specific to regular Dagstuhl Seminars. The schedule was also designed to accommodate the participant's time zones as much as possible. To this end, the schedule was structured around three types of sessions:

- *plenary sessions*: time-zone friendly sessions that constructed the backbone of the seminar (4h per day);
- *Demos & Tech sessions*: sessions reserved for presenting demonstrators and technologies relevant to the seminar (3 sessions of 2h); the objective of these sessions was to ground the discussions and to paint a picture of what can already be achieved with existing technologies;
- *social events*: sessions reserved for online social interactions in a virtual representation of Schloss Dagstuhl via Gather.town[34] (see Figure 6.1).

The seminar started with five invited talks given by James Hendler, Mike Amundsen, Matthias Kovatsch and Simon Mayer (joint talk), Olivier Boissier, and Dave Raggett (in order of presentation). The invited talks were meant to help bootstrap the discussions and presented developments across the three research areas targeted by the seminar. The reduced virtual format did not allow for additional talks from participants, but we organized several rounds of personal introductions during the first two seminar days.

In the third seminar day, participants used a virtual concept board (see Figure 6.2) to organize the seminar topics and to assign them into working groups for the rest of the week. The concept board was created from the seminar topics proposed by the co-organizers, the position statements submitted by participants prior to the seminar, and topics that emerged during the first two seminar days. In total, five working groups were created and four working groups submitted consolidated discussion summaries for the full report (see Section 6 of the full report).

The Demos & Tech sessions attracted more submissions than initially foreseen: we initially scheduled two sessions in the second and third seminar days and eventually scheduled a third session in the fourth seminar day to accommodate all submissions. Submissions were in the form of short abstracts (see Section 5 of the full report) and each participant was given 10 minutes to present a live demonstrator and 5 minutes for questions.

[33] https://purl.org/atac/2021
[34] https://gather.town/, accessed: 14.05.2021.

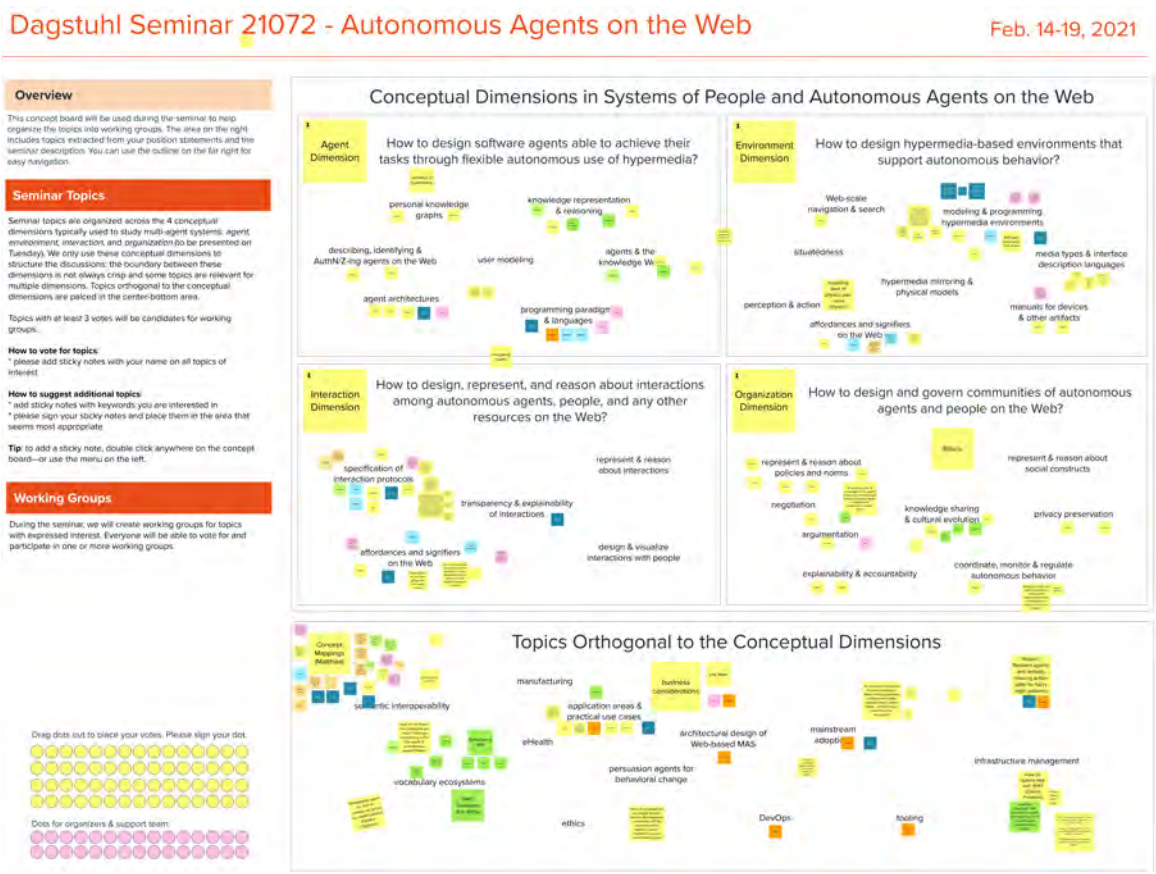

Fig. 6.2
Virtual concept board used to organize the seminar topics into working groups.

Overview of the Report

The full report is organized into four main parts. Section 3 of the full report includes the list of abstracts of the invited talks. Section 4 of the full report includes position statements submitted by participants before and after the seminar. Section 5 of the full report includes the list of abstracts for the demonstrators presented at the seminar. Section 6 of the full report includes the reports submitted by the working groups created during the seminar.

Acknowledgements

The co-organizers would like to thank Fabien Gandon and Simon Mayer for their enthusiasm and support in organizing this seminar. We would also like to thank the team of young researchers who volunteered to support this virtual format: Danai Vachtsevanou (seminar collector), Daniel Schraudner (architect of Virtual Schloss Dagstuhl), Cleber Amaral, Samuele Burattini, Angelo Croatti, Timotheus Kampik, and Adnane Mansour. Finally, we would like to thank Sussanne Bach-Bernhard, Sascha Daeges, and Michael Gerke for their assistance in organizing this seminar in a fully virtual format.

References

1 T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, 284(5):34–43, 2001.

2 J. Hendler. Where are all the intelligent agents? *IEEE Intelligent Systems*, 22(3):2–3, May 2007.

3 J. Hendler and A. Mulvehill. *Social Machines: The Coming Collision of Artificial Intelligence, Social Networking, and Humanity*. Apress, Berkely, CA, USA, 1st edition, 2016.

6.3 Computational Complexity of Discrete Problems

Organizers: Anna Gál, Meena Mahajan, Rahul Santhanam, and Till Tantau

Seminar No. 21121

Date: March 21–26, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.2.1

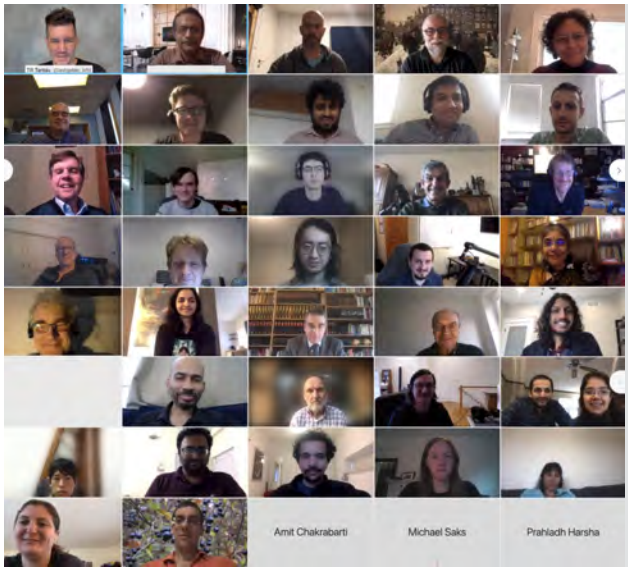

Remote Participants: Eric Allender, Max Bannach, Olaf Beyersdorff, Harry Buhrman, Igor Carboni Oliveira, Amit Chakrabarti, Sourav Chakraborty, Arkadev Chattopadhyay, Lijie Chen, Anindya De, Yuval Filmus, Lance Fortnow, Anna Gál, Mika Göös, Alexander Golovnev, Rohit Gurjar, Kristoffer Arnsfelt Hansen, Prahladh Harsha, Johan Hastad, Shuichi Hirahara, Rahul Ilango, Stacey Jeffery, Gillat Kol, Swastik Kopparty, Michal Koucký, Alexander S. Kulikov, Sophie Laplante, Nutan Limaye, Meena Mahajan, Ian Mertz, Jakob Nordström, Ramamohan Paturi, Toniann Pitassi, Pavel Pudlák, Oded Regev, Rüdiger Reischuk, Robert Robere, Michael E. Saks, Rahul Santhanam, Shubhangi Saraf, Amit Sinhababu, Srikanth Srinivasan, Amnon Ta-Shma, Li-Yang Tan, Till Tantau, Thomas Thierauf, Jacobo Torán, Virginia Vassilevska Williams, Ben Lee Volk, Omri Weinstein, Ryan Williams, Amir Yehudayoff

Computational complexity studies the amount of resources (such as time, space, randomness, or communication) that are necessary to solve computational problems in various models of computation. Finding efficient algorithms for solving computational tasks is crucial for practical applications. Despite a long line of research, for many problems that arise in practice it is not known if they can be solved efficiently – in particular, in polynomial time. Beside questions about the existence of polynomial time algorithms for problems like Satisfiability or Factoring, where the best known algorithms run in exponential time, there is a huge class of practical problems where algorithms with polynomial running time (such as cubic or even quadratic time) are known, but it would be important to establish whether these running times are best possible, to what extent they can be improved, and whether parallel algorithms allow improvements of the runtime.

These fundamental questions motivate developments in various areas from algorithm design to circuit complexity, communication complexity and proof complexity. During this Dagstuhl Seminar 21121, some of the most exciting recent developments in those areas related to computational complexity were presented in a series of talks. The seminar was the most recent one in the series of Dagstuhl Seminars entitled "Computational Complexity of Discrete Problems" – seminars 19121, 17121, 14121, 11121.

Owing to the pandemic and associated travel restrictions, the seminar was held in a purely online format. With 52 researchers from across the world participating in the event, resident in time zones ranging from Japan to California, the window for common acceptable time slots was small. We decided to have a two-hour time slot each day for technical sessions, followed by an additional hour or more each day for social interactions. The Webex platform was used for technical sessions, and gather.town additionally for some of the social interactions. Despite the challenges of making the online event interesting given the ubiquitous screen-time fatigue, the meetings saw high participation level (between at least 80% and typically over 90% participation on all days) and were highly interactive – primarily due to the excellent nature of the given talks.

The seminar started with the creation of a "graph of interests" (using a Miro whiteboard), enabling participants to discover shared research interests with other participants. Following this, during the week, there were 20 research talks, coming from a range of topics including lower bounds on formula size and circuit size, complexity measures of Boolean functions, the algorithmic method for proving lower bounds, fixed parameter tractability and hardness magnification, communication complexity and lifting techniques, as well as proof complexity. Some specific results presented include:

- An improved lower bound, after many years, on the number of hyperplanes needed to slice all edges of the Boolean hypercube.
- A lower bound for monotone arithmetic circuit size using techniques from communication complexity.
- A new potential technique for de Morgan formula lower bounds.
- More refined notions of unambiguous certificate complexity and block sensitivity, with a separation that lifts to communication complexity.

The titles and abstracts of all the talks appear in the full report.

In addition, there was a rump session with short talks by Amit Chakrabarti, Amit Sinhababu, and Prahladh Harsha.

On the social interactions front, in the designated coffee slots

there were some meet-random-people-in-a-break-out sessions. The traditional Wednesday hike was replaced by a "virtual hike" using Google Earth imagery, that went over one of the shorter hike trails near Schloss Dagstuhl and then virtually visited some participants' institutes. The "wine-cheese-music party" became an online party on gather.town following the Schloss Dagstuhl map, and included music, games, and a commentary on the hardness of traveling.

The organizers, Anna Gál, Meena Mahajan, Rahul Santhanam, and Till Tantau, thank all participants for the many contributions they made. We would also like to especially thank the Dagstuhl staff for their cooperation in the current challenging circumstances, their encouragement for going ahead with an online event, and their unstinted help with organizational matters. We would also like to thank Max Bannach for his invaluable help assembling and preparing the report.

6.4 Multi-Level Graph Representation for Big Data Arising in Science Mapping

Organizers: Katy Börner and Stephen G. Kobourov

Seminar No. 21152

Date: April 11–16, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.3.1

Participants: Katy Börner, Ingo Günther, Francis Harvey, Michael Kaufmann, Alexander Wolff

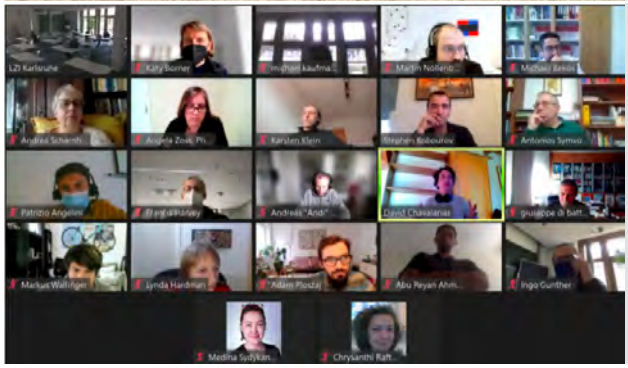

Remote Participants: Patrizio Angelini, Michael A. Bekos, Kevin Boyack, David Chavalarias, Giuseppe Di Battista, Sara Irina Fabrikant, Jean-Daniel Fekete, Lynda Hardman, Ben Jacobsen, Philipp Kindermann, Karsten Klein, Stephen G. Kobourov, Thomas Köhler, Vincent Larivière, Tamara Mchedlidze, Guy Melançon, Staša Milojevic, Filipi Nascimento Silva, Martin Nöllenburg, Adam Ploszaj, Sergey Pupyrev, Chrysanthi Raftopoulou, Andrea Scharnhorst, André Skupin, Cassidy Rose Sugimoto, Antonios Symvonis, Markus Wallinger, Angela Zoss

Networks are all around us. At any moment in time we are driven by, and are an integral part of, many interconnected and dynamically changing networks. Our species has evolved in the context of diverse ecological, biological, social, and other types of networks over thousands of years. We created telephone and power networks, road and airline networks, the Internet and the World Wide Web. We model biological processes with metabolic and protein networks, fake news and rumors with epidemic networks, and even the brain with neural networks. The study of such large-scale networks is one of the prominent generators of big data. Analyzing, exploring, and understanding these complex, interdependent, multi-level networks requires new, more efficient, and more intuitive graph analysis and visualization approaches. This is confirmed in a 2017 VLDB study by Sahu et al. about how graphs are used in practice, where analysts identified scalability and visualization as the most important issues to address.

Recent advances in data, algorithms, and computing infrastructures make it possible to map humankind's collective scholarly knowledge and technology expertise by using topic maps on which "continents" represent major areas of science (e.g., mathematics, physics, or medicine) and zooming reveals successively more detailed subareas. Basemaps of science and technology (S&T) are generated by analyzing citations links between millions of publications and/or patents. "Data overlays" (e.g., showing all publications by one scholar, institution, or country or the career trajectory of a scholar as a pathway) are generated by science-locating relevant publication records based on topical similarity. Science maps are widely used to compare expertise profiles, to understand career trajectories, and to communicate emerging areas. The recent National Academy of Science Colloquium on *Modeling and Visualizing Science Developments* co-organized by Börner showcased the utility of predictive modeling and large-scale mapping efforts, e.g., in support of ranking institutions, analyzing job market developments, and innovation diffusion and technology adoption.

Despite the demonstrated utility of large-scale S&T maps, current approaches do not scale to the hundreds of millions of data records now available. Most users have a hard time reading large-scale networks and few can traverse or derive knowledge from multi-level presentations of networks. Most maps of science, technology, or jobs data support exactly one level of detail. Very few even support two levels such as the UCSD map of science. A key challenge is designing efficient and effective methods

to visualize and interact with more than 100 million scholarly publications at multiple levels of resolution.

Given results from our prior studies on the effectiveness and memorability of map-like visualization of large graphs, we are interested to bring together leading experts to design a multi-level, large-scale map of science that can be used by experts and the general public alike.

The notion of multiple-levels-of-detail graph representation can be captured with Multi-Level Graph Sketches (MLGS) that take a static map-like representation to a multi-level setting needed for exploring and interacting with large, real-world networks. Using the familiar map metaphor, multi-level graph algorithms can make it possible to identify important nodes, major pathways, and clusters across multiple levels. Specifically, we aim to develop efficient algorithms with theoretical and practical guarantees for creating Multi-Level Graph Sketches (MLGS) in support of visual analytics tasks for large network exploration, navigation, and communication. Unlike existing methods for visualizing multi-level networks based on meta-nodes and meta-edges, the MLGS approach can provide real nodes (prototypes) and real paths (backbones) for each level, similar to geographic maps that show real cities and real roads at every level of detail.

Research questions included: (1) designing efficient algorithms for MLGS: advance the state-of-the-art in graph algorithms by generalizing the notion of graph spanners to multiple levels. (2) utilizing MLGS algorithms in visualization: applying the MLGS representation in the context of network analysis and visualization for interacting with large networks, which combines the MLGS approach with clustering, layout, and map-like visualization. (3) developing a new approach for science classification, lookup, and topical mapping service in support of data-driven decision making by students, teachers, and administrators. (4) validating the new algorithms and visualizations: evaluation based on quantitative metrics such as efficiency/scalability, stress/distortion and precision/recall, as well as qualitative metrics based on human subject studies of the utility of visualizations along with readability, engagement, and memorability.

The main goal of this seminar was to bring together researchers coming from information visualization, psychology, cognitive science, human-computer interaction, graph drawing, computational geometry, cartography, and GIS with interests in "science of science" to discuss novel graph mining and layout algorithms and their application to the development of science mapping standards and services.

Due to the pandemic, we had a "hybrid" seminar with only 5 in-person participants and 25 by-zoom participants. With participants from more than 10 different countries and at least 5 different time-zones (some as much as 9 hours behind Central Dagstuhl time), this was a new experience for most of us and definitely different from previous Dagstuhl Seminars. Nevertheless, we attempted to adapt by having additional evening events and moving the traditional Wednesday excursion to the morning.

As this was a highly interdisciplinary seminar, we started the event with talks that introduced the state of the art in the participating fields on a high level. After the introductory presentations, we presented our initial research problems and discussed further questions in an open-problem session.

A set of 4–6 research problems was then finalized and the formation of working groups for these problems was completed by the end of the second day. The remaining three days were dedicated to working-group meetings, progress reports, and initial write-ups.

The main expected outcome of this seminar will be a special issue of the journal *IEEE Computer Graphics and Applications*[35] on the main topics of the seminar. Specifically, we expect 4–6 research papers on the problems discussed by the working groups. Longer term goals include: improved collaborations and communications between the different communities brought together for this seminar, improved maps of science for SciMap2020, and new multi-level graph algorithms and approaches.

[35] https://www.computer.org/digital-library/magazines/cg/call-for-papers-special-issue-on-multi-level-graph-representations-for-big-data-in-science

6.5 Temporal Graphs: Structure, Algorithms, Applications

Organizers: Arnaud Casteigts, Kitty Meeks, George B. Mertzios, and Rolf Niedermeier

Seminar No. 21171

Date: April 25–30, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.3.16

Remote Participants: Evripidis Bampis, Julien Baste, Prithwish Basu, Cristiano Bocci, Hans L. Bodlaender, Binh-Minh Bui-Xuan, Benjamin Bumpus, Chiara Capresi, Arnaud Casteigts, Rémy Cazabet, Andrea Clementi, Timothée Corsini, Pierluigi Crescenzi, Jessica Enright, Thomas Erlebach, Bruno Escoffier, Till Fluschnik, Samuel Hand, Petter Holme, Frank Kammer, Mikko Kivelä, Ralf Klasing, Christian Komusiewicz, Michael Lampis, Zvi Lotker, Andrea Marino, Kitty Meeks, Nicole Megow, George B. Mertzios, Othon Michail, Hendrik Molter, Vincenzo Nicosia, Rolf Niedermeier, William Pettersson, Yoann Pigné, Michael Raskin, Malte Renken, Eric Sanlaville, Jason Schoeters, Fiona Skerman, Martin Skutella, Manuel Sorge, Paul Spirakis, Jakob Spooner, Ondrej Suchý, John Sylvester, Nikolaj Tatti, Amitabh Trehan, Mathilde Vernet, Tiphaine Viard, Victor Zamaraev, Christos Zaroliagis, Philipp Zschoche

Traditionally, graphs (composed of vertices and edges) are used to abstractly model diverse real-world systems, where vertices and edges represent elementary system units and some kind of interactions between them, respectively. However, in modern systems this modeling paradigm using static graphs may be too restrictive or oversimplifying, as interactions often change over time in a highly dynamic manner. The common characteristic in all these application areas is that the system structure, i.e., graph topology, is subject to *discrete changes over time*. In such dynamically changing graphs the notion of *vertex adjacency* needs to be revisited and various graph concepts, e.g., reachability and connectivity, now crucially depend on the exact temporal ordering of the edges' presence. Furthermore, the rate and/or degree of the changes is generally too high to be reasonably modeled in terms of network faults or failures: in these systems changes are not anomalies but rather an integral part of the nature of the system.

A *temporal graph* is a graph which changes over time. Although many different variations for the model of temporal graphs exist, the most common one concerns graphs whose vertex set is fixed while the edge set changes over time. According to this model, given a static underlying graph G, a temporal graph is obtained by assigning to every edge of G a set of nonnegative integer time-labels, indicating the discrete time steps in which this edge is active. Many notions and algorithms on static (i.e., non-temporal) graphs can be transferred in a natural way to their temporal counterpart, while in other cases new approaches are needed to define the appropriate temporal notions. In most cases, the existence of the time dimension adds some degree of freedom in defining a specific temporal variant of a static problem.

The purpose of this one-week seminar was to present and discuss recent advances in the area of temporal and dynamically changing graphs, and to identify and highlight some of the current key challenges to better understand the multiple facets of the computational complexity of temporal graph problems. The seminar was mostly rooted in the rich and mature experience with algorithmic problems on static (i.e., traditional) graphs and in the general quest to understand the computational complexity of temporal versions of fundamental graph problems and algorithms.

Four key areas had been identified, which constituted the backbone of the seminar:

- **Models, Concepts, Classes.** Transferring problems from static to temporal graphs in a meaningful way is a challenging task, with potentially several well-motivated models in the temporal setting corresponding to the same static graph problem. Since temporal graph problems are notoriously hard, a natural way to approach this issue is to restrict input instances. One of the most immediate ways to restrict inputs (which originates in the static algorithmic graph theory) is to restrict the family of underlying graphs G. However, in temporal graphs, the temporal dimension offers so much structure that one can encode hard problems without the need of a sophisticated structure in the underlying graph or in any single layer. This observation suggests that restrictions of the "temporal pattern", i.e., of the way in which the time-labels appear (either alone or in combination with the above mentioned restrictions) can often help to identify tractable special cases of temporal graph problems.
- **Concrete Algorithmic Problems.** There are many examples of canonical generalizations from classic optimization problems on static graphs to temporal graphs. The most meaningful way in which a temporal version of a classic optimization problem may be defined crucially depends on the underlying application domain. Generally, temporal variants tend to become computationally harder than the corresponding classic optimization problem. Canonical ways

to tackle the computational hardness are to (i) restrict the input instances to certain graph classes that allow for efficient (exact) algorithms or find parameters that allow for fixed-parameter algorithms, (ii) aiming for polynomial-time approximation algorithms where a certain level of solution quality is guaranteed, or (iii) combining (i) and (ii).

- **Distributed Aspects.** A common approach to analyzing distributed algorithms is the characterization of necessary and sufficient conditions to their success. These conditions commonly refer to the communication model, synchronicity, or structural properties of the network (e.g., is the topology a tree, a grid, a ring, etc.) In a highly-dynamic network, the topology changes *during* the computation, and this may have a dramatic effect on computation. The study of this impact has led the distributed algorithms community to define a number of classes of temporal graphs that capture various levels of regularities a graph may satisfy over time, regardless of its structure at any single instant. It turns out that some of these properties – in particular the ones pertaining to finite time – are also relevant in a (centralized) algorithmic setting and have an impact on the computational complexity (or feasibility) of classic problems. In particular, properties like periodicity, temporal connectivity, and bounded temporal diameters have been considered both in the distributed and in the non-distributed settings. The seminar played here a crucial role in allowing the different communities to meet and share their complementary experience of temporal properties, as well as to converge on terminology and modeling aspects.
- **Applications.** Whenever there is a situation with pairwise interactions and information about the time point when these interactions happen, the framework of temporal graphs offers a natural mathematical model. This is especially the case in applications where the time information is critical. Examples include traffic and transportation networks, social network analysis (especially when analyzing disease or rumor spreading phenomena), biological networks, mobile sensor networks, and neural networks. All these application areas have their own problem settings and problem instances with specific properties and, in order to apply the fast-developing theory of algorithms for temporal graphs, it is essential to identify and formalise the properties of temporal graphs derived from data sets in these application areas, with the goal of obtaining application-driven restrictions of the input instances that allow for efficient algorithms. Among other application-oriented talks, the seminar included two highly topical talks relating to the role of temporal graphs in pandemic modelling.

The Dagstuhl Seminar 21171 "Temporal Graphs: Structure, Algorithms, Applications" brought together 53 participants from 13 different countries in Europe, USA, Japan, and Israel. The list of participants and speakers contained international experts in algorithms and complexity, social networks and computational social choice, complex systems, distributed algorithms, and parameterized algorithmics. We had in total 21 talks, each of approximately 30 minutes duration, and two sessions where open problems were proposed and discussed.

As this Dagstuhl Seminar was held entirely online (which would be unheard of a few years ago), we provided access to all participants to the online platform GatherTown which enabled us to virtually interact in a way which simulated (very satisfactory) physical meetings and interaction (e.g., also having virtual boards at our disposal). We had the GatherTown platform open every day all day (from the morning until the night), while we specified slots of 2-hours daily where everybody was expected to come to GatherTown. The rest of the day, GatherTown was there to facilitate all people who wanted to have extended physical-looking scientific meetings. During the Seminar, collaborative work was encouraged over all formal and informal scientific discussions. By building on the pre-existing synergies between the participating researchers, new collaborations have been initiated and old ones have been reinforced, and this across all all the communities which were represented in the Seminar.

To conclude, this Seminar has been successful in bringing together scientists from different backgrounds and initiating or strengthening research collaborations on the wide topic of temporal and dynamically changing graphs. Last, but not least, these collaborations and scientific discussions presented a fruitful mix between young researchers, such as PhD students, with older and established researchers in the general field of temporal graphs.

We would like to express our special gratitude to the team of Schloss Dagstuhl who were extremely supportive and also flexible on how to organize a virtual Dagstuhl Seminar during these unprecedented times of the pandemic, without compromising the scientific quality and the overall participation experience of the Seminar.

Acknowledgements The organizers are very grateful to Hendrik Molter and Malte Renken (TU Berlin) for helping in organizing the seminar (program compilation, technical support during the seminar, etc.) and to William Petterson and John Sylvester (University of Glasgow) for their help in collecting the material and compiling the report.

6.6 Computational Geometry

Organizers: Siu-Wing Cheng, Anne Driemel, and Jeff M. Phillips

Seminar No. 21181

Date: May 2–7, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.4.1

Remote Participants: Mikkel Abrahamsen, Pankaj Kumar Agarwal, Helmut Alt, Elena Arseneva, Édouard Bonnet, Karl Bringmann, Mickaël Buchet, Maike Buchin, Sergio Cabello, Hsien-Chih Chang, Siu-Wing Cheng, Man-Kwun Chiu, Jinhee Chun, Éric Colin de Verdière, Jacobus Conradi, Arnaud de Mesmay, Anne Driemel, Ioannis Emiris, Matias Korman, Joseph S. B. Mitchell, Aleksandar Nikolov, Eunin Oh, Steve Y. Oudot, Evanthia Papadopoulou, Zuzana Patáková, Jeff M. Phillips, Benjamin Raichel, Natan Rubin, Maria Saumell, Lena Schlipf, Christiane Schmidt, Donald Sheehy, Kolay Sudeshna, Martin Tancer, Csaba Tóth, Uli Wagner, Carola Wenk, Sue Whitesides

Computational Geometry

The field of 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.

In the early years mostly theoretical foundations of geometric algorithms were laid and fundamental research remains an important issue in the field. Meanwhile, as the field matured, researchers have started paying close attention to applications and implementations of geometric and topological algorithms. Several software libraries for geometric computation (e.g. leda, cgal, core) have been developed. Remarkably, this emphasis on applications and implementations has emerged from the originally theoretically oriented computational geometry community itself, so many researchers are concerned now with theoretical foundations as well as implementations.

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 Computational Topology on Surfaces and Graphs as well as Combinatorial Complexity of Geometric Structures.

Computational Topology on Surfaces and Graphs

Computational topology has seen exciting advances in a number of topics. Indeed, best paper awards in several recent SoCGs went to papers on these topics. In 2019, Cohen-Addad et al. give a lower bound to a cutting problem in embedded graphs, essentially matching the running time of the fastest algorithm known and settling a 17-year old question. In 2018, Goaoc et al. proved that it is NP-complete to decide if a d-dimensional simplicial complex is shellable for $d \geq 2$, resolving a question of Danaraj and Klee in 1978. In 2017, Despré and Lazarus presented simple quasi-linear algorithms for questions regarding geometric intersection number of a curve on a surface. Progress in these and related topics have had influences in problems on graphs embedded on surfaces, maximum flows and multiple-source shortest paths in planar graphs, collapsibility of simplicial complexes, metric learning, etc. The seminar highlighted these topics with two overview talks. The first by Hsien-Chih Chang was on Tightening Curves on Surfaces, and provided a overview of recent advancements in this area, and exciting directions for future work on flipping triangulations and morphing planar multicurves using electrical moves. The second by Uli Wagner discussed Embeddability of Simplicial Compexes, and also the flurry of recent research in this area, and pinpointed the several remaining questions and where the community has not yet been able to resolve the embeddability and why the challenges remain. These talks, and other on recent advances, helped summarize the state of this area, and generate new avenues towards moving the field further forward.

Combinatorial Complexity of Geometric Structures

The understanding of the combinatorial properties of geometric structures is at the core of computational geometry. A lot of these structures such as union of shapes, cuttings, arrangements, Delaunay triangulation, Voronoi diagram have found numerous applications in algorithm design. For example, the analysis of the complexity of the union of translates of a convex body allows

us to understand the complexity of the free space in planning the motion of that convex body under translation. Their studies have also triggered the development of new theoretical tools such as the polynomial method that has been gaining a lot of attention lately. There are also new applications that require the modeling of uncertain data and hence call for a study of many geometric structures under a stochastic setting. The seminar promoted these topics via two overview talks. The first overview talk was by Mikkel Abrahamsen on Minimum Fence Enclosure and Separation Problems; this line of work generalizes the notion of convex hull by identifies other minimally enclosing structures called fences, and the interesting combinatorial properties that arise. The second overview talk by Evanthia Papadopoulou was on Problems in Voronoi and Voronoi-like diagrams. This talk discussed the advancement in generalizations of the classic geometric object of Voronoi diagrams to be defined among geometric objects beyond points, and to higher-order complexes. In addition to providing snapshots of these exciting subareas, they provided future directions for research within these topics and in how they can interact across the broader computational geometry landscape.

Participants and Participation

Dagstuhl Seminars on computational geometry have been organized in a two year rhythm since a start in 1990. They have been extremely successful both in disseminating the knowledge and identifying new research thrusts. Many major results in computational geometry were first presented in Dagstuhl Seminars, and interactions among the participants at these seminars have led to numerous new results in the field. These seminars have also played an important role in bringing researchers together, fostering collaboration, and exposing young talent to the seniors of the field and vice versa. They have arguably been the most influential meetings in the field of computational geometry. The organizers held a lottery for the fifth time this year; the lottery allows to create space to invite younger researchers, rejuvenating the seminar, while keeping a large group of senior and well-known scholars involved. The seminar has now a more balanced attendance in terms of seniority and gender than in the past. This year, 36 researchers from various countries and continents attended the seminar, despite the virtual nature due to COVID-19, showing the strong interest of the community for this event.

Due to the COVID-19 pandemic, the seminar was held entirely virtually. Talks were held over four days. Each day had 2 two-hour blocks of talks, separated by a 2-hour meal break. They were held in the late-afternoon and evening in Europe, which allowed for participants from North America to join in during their morning hours. Unfortunately, this timing was late for those in Asia. The talks were held on Zoom, a Slack server was set up for a more persistent text-based discussion, and a Wonder.me instance was arranged for dynamic forming of group discussions before and after each session. All of these settings were used to communicate research, form collaborations, and attack open problems. Although not as wonderful as actually being at Schloss Dagstuhl, these online mechanisms provided for a workable replacement for what a normal Dagstuhl Seminar provides in this abnormal time.

The feedback from participants was very positive. The participants viewed the composition of the group positively, remarking how it was well-balanced in terms of seniority and gender. They also praised the quality of the talks as of very high quality – making the virtual-only participation worthwhile.

We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl made virtual hosting possible and easy in a time filled with complications. Despite not providing a physical space to meet, socialize, and collaborate, their help in organizing the event made it a success despite the less than ideal circumstances.

6.7 Approaches and Applications of Inductive Programming

Organizers: Andrew Cropper, Luc De Raedt, Richard Evans, and Ute Schmid

Seminar No. 21192

Date: May 9–12, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.4.20

Remote Participants: Lun Ai, Nada Amin, Martin Atzmüller, Feryal Behbahani, Thea Behrens, Andrew Cropper, James Cussens, Wang-Zhou Dai, Luc De Raedt, Thomas Demeester, Amit Dhurandhar, Sebastijan Dumancic, Kevin Ellis, Richard Evans, Cesar Ferri Ramirez, Bettina Finzel, Peter Flach, Johannes Fürnkranz, Artur Garcez, Manuel Garcia-Piqueras, José Hernández-Orallo, Céline Hocquette, Frank Jäkel, Gonzalo Jaimovitch López, Susumu Katayama, Tomáš Kliegr, Stefan Kramer, Maithilee Kunda, Sara Magliacane, Roman Manevich, Fernando Martinez-Plumed, Pasquale Minervini, Stephen H. Muggleton, Stassa Patsantzis, Johannes Rabold, Claude Sammut, Stephan Scheele, Ute Schmid, Javier Segovia-Aguas, Gustavo Soares, Stefano Teso, Jan Tinapp

The goal of Inductive Programming (IP) is to provide methods for induction of computer programs from data. Specifically, IP is the automated (or semi-automated) generation of a computer program from an incomplete information, such as input-output examples, demonstrations, or computation traces.IP offers powerful approaches to learning from relational data and to learning from observations in the context of autonomous intelligent agents. IP is a form of machine learning, because an IP system should perform better given more data (i.e. more examples or experience). However, in contrast to standard ML approaches, IP approaches typically only need a small number of training examples. Furthermore, induced hypotheses are typically represented as logic or functional programs, and can therefore be inspected by a human. In that sense, IP is a type of interpretable machine learning which goes beyond the expressivity of other approaches of rule learning such as decision tree algorithms. IP is also a form of program synthesis. It complements deductive and transformational approaches. When specific algorithm details are difficult to determine, IP can be used to generate candidate programs from either user-provided data, such as test cases, or from data automatically derived from a formal specification. Most relevant application areas of IP techniques is end-user programming and data wrangling.

This seminar has been the fifth in a series – building on seminars 13502, 15442, 17383, and 19202. In the wake of the recent interest in deep learning approaches, mostly for end-to-end learning, it has been recognized that for practical applications, especially in critical domains, data-intensive blackbox machine learning must be complemented with methods which can help to overcome problems with data quality, missing or errouneous labeling of training data, as well as providing transparency and comprehensibility of learned models. To address these requirements, on the one hand, explainable artificial intelligence (XAI) emerged as a new area of research and on the other hand, there is a new interest in bringing together learning and reasoning. These two areas of research are in the focus of the 2021 seminar. Futhermore, recent developments to scale up IP methods to be more applicable to complex real world domains has been taken into account. Based on outcomes of the fourth seminar (19202), the potential of IP as powerful approach for explainable artificial intelligence ("IP for XAI") has been be elaborated. Bringing together IP methods and deep learning approaches contributes to neural-symbolic intergration research. While two years ago (seminar 19202) focus has been on IP as interpretable surrogate model, in the 2021 seminar explainability of different addressees of explanations and their need to different types of explanations (e.g. verbal or example-based) are considered. For many real world applications, it is necessary to involve the human as teacher and judge for the machine learned models. Therefore, a further topic of the seminar has been to explore IP in the context of new approaches to interactive ML and their applications to automating data science and joint human-computer decision making.

6.8 Serverless Computing

Organizers: Cristina Abad, Ian T. Foster, Nikolas Herbst, and Alexandru Iosup

Seminar No. 21201

Date: May 16–21, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.4.34

Participants: André Bauer, Simon Eismann, Nikolas Herbst

Remote Participants: Cristina Abad, Samer Al-Kiswany, Ahmed Ali-Eldin Hassan, Bartosz Balis, André B. Bondi, Kyle Chard, Ryan L. Chard, Robert Chatley, Andrew A. Chien, A. Jesse Jiryu Davis, Jesse Donkervliet, Erik Elmroth, Nicola Ferrier, Ian T. Foster, Alexandru Iosup, Hans-Arno Jacobsen, Pooyan Jamshidi, Samuel Kounev, Georgios Kousiouris, Philipp Leitner, Pedro García López, Martina Maggio, Maciej Malawski, Bernard Metzler, Vinod Muthusamy, Alessandro Vittorio Papadopoulos, Panos Patros, Guillaume Pierre, Omer F. Rana, Robert P. Ricci, Joel Scheuner, Mina Sedaghat, Mohammad Shahrad, Prashant Shenoy, Josef Spillner, Davide Taibi, Douglas Thain, Animesh Trivedi, Alexandru Uta, Vincent van Beek, Erwin van Eyk, André van Hoorn, Soam Vasani, Florian Wamser, Guido Wirtz, Vladimir Yussupov

Serverless computing holds a significant promise for the modern, digital society. For the past seven decades, our society has increasingly required ever-cheaper, ever-more convenient, and ever-faster computing technology. In the late-1950s, leasing time on an IBM 701 cost $15,000 per month ($135,000 in 2020 dollars). Today, we can lease many times this computing power for mere pennies but need to be careful about the actual cost of doing so. Cloud computing, that is, the utility providing IT as a service, on-demand and pay-per-use, is a widely used computing paradigm that offers large economies of scale and promises extreme environmental efficiency. Born from a need to make cloud computing services more accessible, fine-grained, and affordable, *serverless computing* has garnered interest from both industry and academia. In our vision, serverless computing can meet this need, but to do this it will have to overcome its current status of emergent technology or risk its demise.

Cloud computing is already an established technology. Today, more than three-quarters of the US and European companies, and many private individuals, use cloud computing services[36]. The serverless market is blooming[37] and has already exceeded $200 billion in 2020[38]. The cost of one hour on a cloud computer leased on-demand can be lower than a cent[39] and all the major cloud providers offer inexpensive access to diverse and state-of-the-art hardware. However cheap, cloud computing still poses daunting

[36] European Commission, Uptake of Cloud in Europe, Digital Agenda for Europe report by the Publications Office of the European Union, Luxembourg, Sep 2014. and Flexera, State of the Cloud Report, 2020.
[37] Gartner Inc. Gartner Forecasts Worldwide Public Cloud Revenue to Grow 17% in 2020. Press Release.
[38] Frank Gens. Worldwide and Regional Public IT Cloud Services 2019–2023 Forecast. Tech. Rep. by IDC, Doc. #US44202119, Aug 2019.
[39] Amazon AWS, Microsoft Azure, and Google Compute Engine offer VMs in this price range.

operational challenges to software professionals, in particular, how to manage the selection, operation, and other aspects of using cloud infrastructure (in short, *servers*). Correspondingly, it poses significant challenges to systems designers and administrators, related to keeping the cloud infrastructure efficient and sustainable.

An emerging class of cloud-based software architectures, *serverless computing*, focuses on providing software professionals the ability to execute arbitrary functions with low or even no overhead in server management. Serverless computing leverages recent developments in the miniaturization of software parts through *microservice-based architectures*, in the operationalization of small self-contained execution units through *containers*, and in their integration in service models such as *Function-as-a-Service (FaaS)*. Truly, serverless is more [12]. Early research successes [6, 15, 17, 18, 22] complement numerous industrial applications [9], from business-critical to scientific computing, from DevOps to side-tasks. Already, IT spending on serverless computing should exceed $8 billion per year, by 2021.[40]

However promising, serverless computing has yet to mature and presents many hard, open challenges. There are numerous signs and reports [11, 14] that serverless computing poses critical challenges in software engineering, parallel and distributed systems operation, and performance engineering [10]. For example, software engineering could help overcome challenges in the developer experience [23], including testing, tooling, functionality, and training and education. The systems side requires, among others, new approaches for deployment, monitoring, and general operation, and also specific advances in security, cost predictability, and life-cycle management for cloud functions. Performance engineering raises many hard aspects, such as performance optimization, engineering for cost-efficiency, and various forms of fast online scheduling. These combined challenges are distinctive from the general challenges of cloud computing, for example, because the fine-grained, often event-driven nature of serverless computing typically requires approaches that are lightweight and able to respond without delay.

The goal of the seminar is to *combine the views of a diverse and high-quality group of researchers spanning three disciplines*: software engineering, parallel and distributed systems, and performance engineering. The Dagstuhl Seminar will be a catalyst. Attendees discussed the open challenges and opportunities of serverless computing for the next decade, with a focus on at least the following crucial aspects and questions:

- Envision serverless systems and applications in the next decade. How to leverage the freedom from operational concerns? How to overcome the challenge and enjoy the benefits of fine granularity?
- How to properly engineer serverless software and systems? What are the emerging architectural patterns for serverless systems and applications? How to test and debug serverless systems and applications?
- How to characterize, model, and analyze serverless systems and applications? How to understand the diverse serverless workloads?
- How to manage the resources used in serverless operations? How to schedule and orchestrate in this environment? How to manage specific application classes, such as computer vision, enterprise workflows, HPC, DevOps?
- How to deploy and manage the full lifecycle of serverless applications? How to add ML-capabilities to feedback loops? How to break through the operational silos?
- How to support privacy, security, dependability, and other desirable operational properties for serverless applications and systems?
- Beyond computer systems, how to consider serverless systems and applications from a holistic, cyberphysical perspective?

Core topics

The seminar focussed on the following key topics related to serverless computing:

Topic 1. Design decisions for serverless systems, platforms, and ecosystems. As the serverless model is increasingly being adopted in industry [9], the challenges of properly designing these systems and the platforms on which they run are becoming more apparent. These challenges include important problems [10], such as: how to reduce the serverless overhead added by the platform to the (commonly lightweight) functions representing the business logic of the application (e.g., see [20]), how to ensure proper performance isolation while making efficient use of the shared infrastructure (e.g., see [1]), how to partition the functions [5, 6], and how to properly schedule functions and route requests to these functions (e.g., see [2]), in such a way that the service level objectives (SLO's) are adequately met, among other important challenges. There is also the question of running serverless workloads alongside conventional applications, e.g., HPC, big data, machine learning. The experiences of the attendees to the seminar, some of which have already started working in these domain and others with established experience in prior technologies from which we may learn and transfer knowledge (e.g., grid computing), will enable us to focus on determining which of these decisions the community should be focusing on, and how to establish adequately prioritized research agendas.

Topic 2. Software engineering of serverless applications, but also systems, platforms, and ecosystems. To increase the domain of application for serverless computing, the functionality it can express needs to become increasingly more complex, which contrasts with the perceived simplicity of the model [23]. What is the trade-off between simplicity and expressiveness? Which composition models can ensure that serverless workflows can be maintained and developed (and updated) long term? Serverless functions should become increasingly interoperable, and applications should become able to leverage the services of any serverless platform [6]. How to make serverless functions vendor-agnostic and how to run serverless applications across cloud federations? Which architectural patterns are useful for serverless applications? How to consider and support the legacy part of serverless applications? The development processes, from the macro-view of how teams coordinate to deliver applications that could operate in an external ecosystem, to the micro-view of how to develop and test a serverless function, will have to consider the new aspects raised by serverless computing. What are effective development processes? What tools and IDE features are needed? What versioning and testing, and what CI/CD protocols should be used? How to evolve legacy software toward serverless-native applications? How to ensure open-source software becomes FAIR software [13]?

Topic 3. Applications and domain requirements for serverless computing. Preliminary studies of serverless applications at large [9] have shown that there is a wide variety of

[40] "Function-as-a-Service Market – Global Forecast to 2021," marketsandmarkets.com, Feb 2017.

scenarios for which industry and academia are adopting serverless approaches. From business-critical workloads, to automating DevOps, scientific computing, and beyond, the diversity of the applications and domains for which serverless is being applied poses significant challenges when attempting to optimally manage the resources and infrastructure on which these applications depend. It is important to properly understand the variety of these applications and domain requirements, engaging both academia and industry in the discussion.

These requirements should relate to various aspects in software engineering, parallel and distributed systems, and performance engineering. For example, a domain-based approach could help increase scalability [3]; considering the structure of packages in composing a deployable serverless application could improve scheduling performance [2]; and serverless functions and architectures should be considered during performance tests [8, 28].

Topic 4. Evaluation of serverless computing systems, platforms, and ecosystems. The performance trade-offs of serverless systems are not yet well understood [28], thus highlighting the importance of proper evaluation and benchmarking of these systems. However, the high level of abstraction and the opaqueness of the operational-side make evaluating these platforms particularly challenging. As recent efforts are starting to focus on this topic [24, 28], it is important to engage the community on an early discussion on the best approaches to tackle this problem. How to understand and engineer the performance of serverless systems? How to translate the findings, when serverless systems are opened to external developers (as platforms) or take part in much larger systems of systems (and even ecosystems)? How to account for parts of the ecosystem being closed-source and even acting as black-boxes? How to identify and even explain the performance bottlenecks such systems, platforms, and ecosystems experience? How to use evaluation results with other performance engineering techniques to control and improve the performance of serverless systems, platforms, and ecosystems?

An important focus of inquiry has recently become prominent in computer systems research: the reproducibility of evaluation results and of experiments in general [19, 21]. Not doing so can result in misleading results [26], and in results that cannot be obtained again [25] sometimes even under identical circumstances and by their original authors [7]. This leads to a possible loss of faith in the entire field [4, 27]. "How to benchmark serverless solutions reproducibly?" is an important question to address with diverse expertise and fresh ideas.

Synopsis and Planned Actions We would like to thank the Dagstuhl staff and sponsors for this unique seminar opportunity even under the constraints of the pandemic. During the seminar, we had almost 45h of online meetings (not counting sub-meetings): some 9-10h of online meetings each seminar day. Three 3h sessions per day were spread around the clock to allow participation from various timezones. Even under these constraints, we experienced enormous participation and active discussion involvement. In brief, the seminar week was structured as follows:

After each participant presented her/himself to the plenary, we formed four working groups according to the topics above. The discussions were kick-started by four distinguished keynotes, in plenary, with the respective talk abstracts included in the full report:

- "Serverless Predictions: 2021-2030" given jointly by Pedro García López (Universitat Rovira i Virgili – Tarragona, ES) and Bernard Metzler (IBM Research-Zurich, CH)
- "Developer Experience for Serverless: Challenges and Opportunities" given by Robert Chatley (Imperial College London, GB)
- "Federated Function as a Service" given jointly by Kyle Chard (University of Chicago, US) and Ian T. Foster (Argonne National Laboratory – Lemont, US)
- "Characterizing Serverless Systems" given by Mohammad Shahrad (University of British Columbia – Vancouver, CA)

Each of the four working groups held five 3h sessions with their teams, including three 1h one-on-one meetings with the other groups. The four working groups report individually on their outcomes and list identified research challenges. In a consolidation phase, we identified and planned nine focused topics for future joint research among the participants.

Complemented by a Slack workspace for the seminar participants, a focused continuation of discussions beyond the seminar week was enabled: Among others, a discussion initiated and led by Samuel Kounev on the notion of serverless computing, started during the seminar, continued well beyond. We include the outcome of this "panel discussion" in Section 5.1 of the full report.

The organizers and participants decided to jointly work toward at least one high-profile magazine article reporting on the seminar outcome and research agenda.

Furthermore, during the seminar the motion was raised to establish a conference series on serverless computing. We see good potential for a new conference on "Serverless Software and Systems" as a cross-community event embracing, at least, the disciplines of software engineering, system engineering, and performance engineering. Working potentially in concert with an existing workshop series in the field, we plan to initiate this step in the coming months. We hope that one day in the future, we can proudly look back and say that this Dagstuhl Seminar 21201 was an important trigger event.

References

1 Z. Al-Ali, S. Goodarzy, E. Hunter, S. Ha, R. Han, E. Keller, and E. Rozner. Making serverless computing more serverless. In 2018 IEEE 11th International Conference on Cloud Computing (CLOUD), 2018.

2 G. Aumala, E. Boza, L. Ortiz-Avilés, G. Totoy, and C. Abad. Beyond load balancing: Package-aware scheduling for serverless platforms. In 2019 19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), 2019.

3 Alberto Avritzer, Vincenzo Ferme, Andrea Janes, Barbara Russo, André van Hoorn, Henning Schulz, Daniel S. Menasché, and Vilc Queupe Rufino. Scalability assessment of microservice architecture deployment configurations: A domain-based approach leveraging operational profiles and load tests. J. Syst. Softw., 165:110564, 2020.

4 Monya Baker. Is there a reproducibility crisis? Nature, 533(7604):452–454, 2016.

5 Edwin F. Boza, Xavier Andrade, Jorge Cedeno, Jorge R. Murillo, Harold Aragon, Cristina L. Abad, and Andres G. Abad. On implementing autonomic systems with a serverless computing approach: The case of self-partitioning cloud caches. Comput., 9(1):14, 2020.

6 Ryan Chard, Yadu N. Babuji, Zhuozhao Li, Tyler J. Skluzacek, Anna Woodard, Ben Blaiszik, Ian T. Foster, and Kyle Chard. funcx: A federated function serving fabric for science. In HPDC '20: The 29th International Symposium on High-Performance Parallel and Distributed Computing, pages 65–76. ACM, 2020.

7 Christian S. Collberg and Todd A. Proebsting. Repeatability in computer systems research. Commun. ACM, 59(3):62–69, 2016.

8 Simon Eismann, Cor-Paul Bezemer, Weiyi Shang, Dusan Okanovic, and André van Hoorn. Microservices: A performance tester's dream or nightmare? In ICPE '20: ACM/SPEC International Conference on Performance Engineering, 2020, pages 138–149. ACM, 2020.

9 Simon Eismann, Joel Scheuner, Erwin van Eyk, Maximilian Schwinger, Johannes Grohmann, Nikolas Herbst, Cristina L. Abad, and Alexandru Iosup. A review of serverless use cases and their characteristics. Technical Report SPEC-RG-2020-5, SPEC RG Cloud Working Group, May 2020.

10 Erwin Van Eyk, Alexandru Iosup, Cristina L. Abad, Johannes Grohmann, and Simon Eismann. A SPEC RG cloud group's vision on the performance challenges of FaaS cloud architectures. In Companion of the 2018 ACM/SPEC International Conference on Performance Engineering, ICPE 2018, pages 21–24, 2018.

11 Erwin Van Eyk, Alexandru Iosup, Simon Seif, and Markus Thömmes. The SPEC cloud group's research vision on FaaS and serverless architectures. In Proceedings of the 2nd International Workshop on Serverless Computing, WOSC@Middleware 2017, Las Vegas, NV, USA, December 12, 2017, pages 1–4, 2017.

12 Erwin Van Eyk, Lucian Toader, Sacheendra Talluri, Laurens Versluis, Alexandru Uta, and Alexandru Iosup. Serverless is more: From PaaS to present cloud computing. IEEE Internet Comput., 22(5):8–17, 2018.

13 Wilhelm Hasselbring, Leslie Carr, Simon Hettrick, Heather S. Packer, and Thanassis Tiropanis. From FAIR research data toward FAIR and open research software. it Inf. Technol., 62(1):39–47, 2020.

14 Joseph M. Hellerstein, Jose M. Faleiro, Joseph Gonzalez, Johann Schleier-Smith, Vikram Sreekanti, Alexey Tumanov, and Chenggang Wu. Serverless computing: One step forward, two steps back. In CIDR 2019, 9th Biennial Conference on Innovative Data Systems Research, Asilomar, CA, USA, January 13-16, 2019, Online Proceedings. www.cidrdb.org, 2019.

15 Scott Hendrickson, Stephen Sturdevant, Tyler Harter, Venkateshwaran Venkataramani, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. Serverless computation with openlambda. In 8th USENIX Workshop on Hot Topics in Cloud Computing, HotCloud 2016, Denver, CO, USA, June 20-21, 2016., 2016.

16 Alexandru Iosup, Catia Trubiani, Anne Koziolek, José Nelson Amaral, Andre B. Bondi, and Andreas Brunnert. Flexibility is key in organizing a global professional conference online: The ICPE 2020 experience in the COVID-19 era. CoRR, abs/2005.09085, 2020.

17 Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-che Tsai, Anurag Khandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Jayant Yadwadkar, Joseph E. Gonzalez, Raluca Ada Popa, Ion Stoica, and David A. Patterson. Cloud programming simplified: A berkeley view on serverless computing. CoRR, abs/1902.03383, 2019.

18 Ana Klimovic, YawenWang, Patrick Stuedi, Animesh Trivedi, Jonas Pfefferle, and Christos Kozyrakis. Pocket: Elastic ephemeral storage for serverless analytics. In 13th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2018, Carlsbad, CA, USA, October 8-10, 2018, pages 427–444. USENIX Association, 2018.

19 Ravi Madduri, Kyle Chard, Mike D'Arcy, Segun C. Jung, Alexis Rodriguez, Dinanath Sulakhe, Eric Deutsch, Cory Funk, Ben Heavner, Matthew Richards, Paul Shannon, Gustavo Glusman, Nathan Price, Carl Kesselman, and Ian Foster. Reproducible big data science: A case study in continuous fairness. PLOS ONE, 14(4):1–22, 04 2019.

20 Edward Oakes, Leon Yang, Dennis Zhou, Kevin Houck, Tyler Harter, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau. SOCK: Rapid task provisioning with serverlessoptimized containers. In USENIX Annual Technical Conference (USENIX ATC 18), July 2018.

21 A. V. Papadopoulos, L. Versluis, A. Bauer, N. Herbst, J. Von Kistowski, A. Ali-eldin, C. Abad, J. N. Amaral, P. Tuma, and A. Iosup. Methodological principles for reproducible performance evaluation in cloud computing. IEEE Trans. Software Eng., 2019.

22 Qifan Pu, Shivaram Venkataraman, and Ion Stoica. Shuffling, fast and slow: Scalable analytics on serverless infrastructure. In Jay R. Lorch and Minlan Yu, editors, 16th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2019, Boston, MA, February 26-28, 2019, pages 193–206. USENIX Association, 2019.

23 Mike Roberts. Serverless architectures. https://martinfowler.com/articles/serverless.html, 2016. Continuous development of the material.

24 Joel Scheuner and Philipp Leitner. Function-as-a-service performance evaluation: A

multivocal literature review. Journal of Systems and Software, 2020.
25 Dag I. K. Sjøberg, Jo Erskine Hannay, Ove Hansen, Vigdis By Kampenes, Amela Karahasanovic, Nils-Kristian Liborg, and Anette C. Rekdal. A survey of controlled experiments in software engineering. IEEE Trans. Software Eng., 31(9):733–753, 2005.
26 Alexandru Uta, Alexandru Custura, Dmitry Duplyakin, Ivo Jimenez, Jan S. Rellermeyer, Carlos Maltzahn, Robert Ricci, and Alexandru Iosup. Is big data performance reproducible in modern cloud networks? In 17th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2020, Santa Clara, CA, USA, February 25-27, 2020, pages 513–527, 2020.
27 Erik van der Kouwe, Gernot Heiser, Dennis Andriesse, Herbert Bos, and Cristiano Giuffrida. Benchmarking flaws undermine security research. IEEE Secur. Priv., 18(3):48–57, 2020.
28 Erwin van Eyk, Joel Scheuner, Simon Eismann, Cristina L. Abad, and Alexandru Iosup. Beyond microbenchmarks: The SPEC-RG vision for a comprehensive serverless benchmark. In Companion of the ACM/SPEC International Conference on Performance Engineering (ICPE), page 26–31, 2020.

6.9 Transparency by Design

Organizers: Judy Kay, Tsvi Kuflik, and Michael Rovatsos

Seminar No. 21231

Date: June 6–10, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.5.1

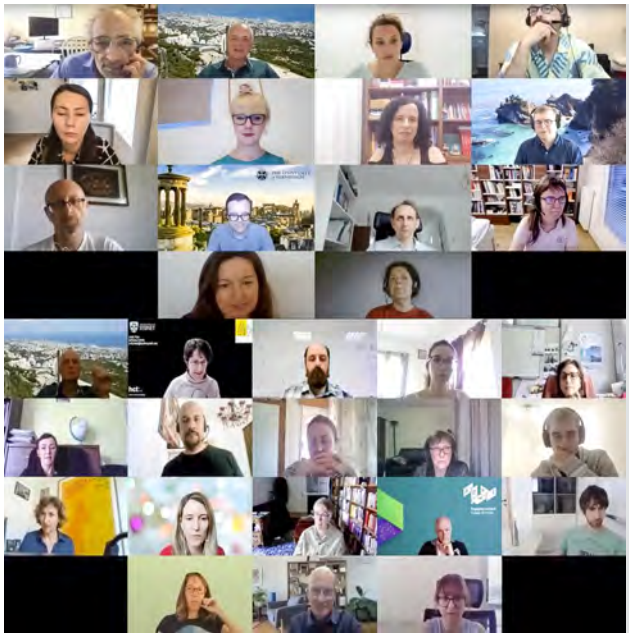

Remote Participants: Elisabeth André, Bettina Berendt, Nehal Bhuta, Maria Bielikova, Veronika Bogina, Joanna J. Bryson, Robin Burke, Aylin Caliskan, Carlos Castillo, Ewa Cepil, Cristina Conati, Aviva de Groot, Gianluca Demartini, Virginia Dignum, Fausto Giunchiglia, Riccardo Guidotti, Meeri Haataja, Judy Kay, Styliani Kleanthous, Ansgar Koene, Joshua A. Kroll, Antonio Krüger, Tsvi Kuflik, Bob Kummerfeld, Loizos Michael, Antonija Mitrovic, Kalia Orphanou, Jahna Otterbacher, Anna Perini, Lena Podoletz, Iris Reinhartz-Berger, Paolo Rosso, Michael Rovatsos, Avital Shulner-Tal, Alison Smith-Renner, Andreas Theodorou, Vincent Wade, Emine Yilmaz

As AI technologies are witnessing impressive advances and becoming increasingly widely adopted in real-world domains, the debate around the ethical implications of AI has gained significant momentum over the last few years. Much of this debate has focused on fairness, accountability, transparency and ethics, giving rise to "Fairness, Accountability and Transparency" (FAT or FAccT) being commonly used to capture this complex of properties as key elements to ethical AI.

However, the notion of transparency – closely linked to terms like explainability, accountability, and interpretability – has not yet been given a holistic treatment within computer science. Despite the fact that it is a prerequisite to instilling trust in AI technologies, there is a gap in understanding around how to create systems with the required transparency, from demands on capturing their transparency requirements all the way through to concrete design and implementation methodologies. When it comes to, for example, demonstrating that a system is fair or accountable, we lack usable theoretical frameworks for transparency. More generally, there are no general practical methodologies for the design of transparent systems.

The purpose of this Dagstuhl Seminar was to initiate a debate around theoretical foundations and practical methodologies with the overall aim of laying the foundations for a "Transparency by Design" framework, i.e. a framework for systems development that integrates transparency in all stages of the software development process.

To address this challenge, we brought together researchers with expertise in Artificial Intelligence, Human-Computer Interaction, and Software Engineering, but also considered it essential to invite experts from the humanities, law and social sciences, which would bring an interdisciplinary dimension to the seminar to investigate the cognitive, social, and legal aspects of transparency.

As a consequence of the Covid-19 pandemic, the seminar had to be carried out in a virtual, online format. To accommodate the time zones of participants from different parts of the world, two three-hour sessions were scheduled each day, with participant groups of roughly equal size re-shuffled each day to provide every attendee with opportunities to interact with all other participants whenever time difference between their locations made this possible in principle. Each session consisted of plenary talks and discussion as well as work in small groups, with discussions and outcomes captured in shared documents that were edited jointly by the groups attending different sessions each day.

The seminar was planned to gradually progress from building a shared understanding of the problem space among participants on the first day, to mapping out the state of the art and identifying gaps in their respective areas of expertise on the second day and third day.

To do this, the groups identified questions that stakeholders in different domains may need to be able to answer in a transparent systems, where we relied on participants to choose domains they are familiar with and consider important. To identify the state of the art in these areas, the group sessions on the second and third days were devoted to mapping out the current practice and research, identifying gaps that need to be addressed.

The two sessions on each day considered these in terms of

four aspects: data collection techniques, software development methodologies, AI techniques and user interfaces.

Finally, the last day was dedicated to consolidating the results towards creating a framework for designing transparent systems. This began with each of the parallel groups considering different aspects: Motivating why transparency is important; challenges posed by current algorithmic systems; transparency-enhancing technologies; a transparency by design methodology; and, finally, the road ahead.

The work that began with the small group discussions and summaries continued with follow up meetings to continue the work of each group. The organisers have led the work to integrate all of these into an ongoing effort after the seminar, aiming to create a future joint publication.

6.10 Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments

Organizers: Stephen Brewster, Andrew Kun, Andreas Riener, and Orit Shaer

Seminar No. 21232

Date: June 6–11, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.5.23

Remote Participants: Sun Joo Ahn, Ignacio J. Alvarez, Laura Boffi, Susanne Boll, Stephen Brewster, Duncan Brumby, Gary Burnett, Yi-Chao Chen, Lewis Chuang, Anna Cox, Birsen Donmez, Geraldine Fitzpatrick, Mark D. Gross, Shamsi Tamara Iqbal, Christian P. Janssen, Wendy Ju, Andrew Kun, John D. Lee, Siân Lindley, Mark McGill, Helena M. Mentis, Alexander Meschtscherjakov, Bastian Pfleging, Champika Manel Epa Ranasinghe, Andreas Riener, Sayan Sarcar, Clemens Schartmüller, Martina Schuß, Orit Shaer, Gregory F. Welch

Agenda in a nutshell

The seminar was conducted online during the week of June 6-10, 2021. A particular difficulty in planning the agenda (see Fig. 6.3) arose due to the different time zones of the individual participants. It was especially important to us to offer at least some of the program items together to all participants (Opening and Group Work on Day 1, Summit and Closing on the last day). On the other days, we planned different activities in smaller (2-3 people) and larger (up to half of the participants) groups to better accommodate the participants based on their time zones.

Dagstuhl Seminar 21232 "Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments"

					all	all	individual/small groups	individual/small groups	all
Europe (CEST, Germany)	US (EDT, e.g. Boston MA)	US (West, Seattle SA)	China/Hong Kong	Tokyo, Japan / Australia	Monday, June 7	Tuesday, June 8	Wednesday, June 9	Thursday, June 10	Friday, June 10
12:00:00 am	6:00:00 pm	3:00:00 pm	6:00:00 am	7:00:00 am					
1:00:00 am	7:00:00 pm	4:00:00 pm	7:00:00 am	8:00:00 am				Video discussions (pairs)	
2:00:00 am	8:00:00 pm	5:00:00 pm	8:00:00 am	9:00:00 am				Compilation of playlists (pairs)	
3:00:00 am	9:00:00 pm	6:00:00 pm	9:00:00 am	10:00:00 am					
4:00:00 am	10:00:00 pm	7:00:00 pm	10:00:00 am	11:00:00 am					
5:00:00 am	11:00:00 pm	8:00:00 pm	11:00:00 am	12:00:00 pm					
6:00:00 am	12:00:00 am	9:00:00 pm	12:00:00 pm	1:00:00 pm					
7:00:00 am	1:00:00 am	10:00:00 pm	1:00:00 pm	2:00:00 pm					
8:00:00 am	2:00:00 am	11:00:00 pm	2:00:00 pm	3:00:00 pm					
9:00:00 am	3:00:00 am	12:00:00 am	3:00:00 pm	4:00:00 pm		Workshop1 G1 (19)			
10:00:00 am	4:00:00 am	1:00:00 am	4:00:00 pm	5:00:00 pm		Workshop1 G1	Movie "Coded bias" +discussion (Miro)		
11:00:00 am	5:00:00 am	2:00:00 am	5:00:00 pm	6:00:00 pm		Workshop1 G1	Movie "Coded bias" +discussion (Miro)		
12:00:00 pm	**6:00:00 am**	**3:00:00 am**	**6:00:00 pm**	**7:00:00 pm**					
1:00:00 pm	7:00:00 am	4:00:00 am	7:00:00 pm	8:00:00 pm					
2:00:00 pm	8:00:00 am	5:00:00 am	8:00:00 pm	9:00:00 pm					
3:00:00 pm	9:00:00 am	6:00:00 am	9:00:00 pm	10:00:00 pm					
4:00:00 pm	10:00:00 am	7:00:00 am	10:00:00 pm	11:00:00 pm					
5:00:00 pm	11:00:00 pm	8:00:00 am	11:00:00 pm	12:00:00 am	Opening & Introduction	Workshop2 G2 (US, 11)			Summit
6:00:00 pm	12:00:00 pm	9:00:00 am	12:00:00 am	1:00:00 am	Opening & Introduction	Workshop2 G2 (US)	Movie "Coded bias" +discussion (Miro)		Summit
7:00:00 pm	1:00:00 pm	10:00:00 am	1:00:00 am	2:00:00 am	Opening & Introduction	Workshop2 G2 (US)	Movie "Coded bias" +discussion (Miro)		Summit
8:00:00 pm	2:00:00 pm	11:00:00 am	2:00:00 am	3:00:00 am					
9:00:00 pm	3:00:00 pm	12:00:00 pm	3:00:00 am	4:00:00 am					
10:00:00 pm	4:00:00 pm	1:00:00 pm	4:00:00 am	5:00:00 am					
11:00:00 pm	5:00:00 pm	2:00:00 pm	5:00:00 am	6:00:00 am					

Fig. 6.3
Compact overview of the agenda for the week including different geographical zones (for better planning with participants from all-over the world).

- Monday, June 6: The seminar was opened and its main goals introduced by the seminar co-organizers Stephen Brewster, Andrew Kun, Andreas Riener and Orit Shaer. The presented slides can be accessed here: https://docs.google.com/presentation/d/15NtQy96wAS_dMHpdqTO-2TfZhbSGssWAn96RxHJizVA/edit#slide=id.gdd9402fdbb_0_8. After a social "warm-up" activity, Pecha Kucha presentations of all participants followed. During the presentations, all participants were instructed to collect questions, ideas, thoughts, etc. on a Miro-board; The items were clustered by the organizers (in a short coffee break) and after that, a voting of topics to be picked-up/focusing on in the next days of the seminar (see Fig. 6.4) followed. This activity ended day 1.
- Tuesday, June 7: The second day of the seminar was dedicated to the "Work(shop) for the Future of Work and Mobility in Automated Vehicles". In this workshop, participants (see Fig. 6.5) worked together on user needs and how to fulfill them during shared or private automated mobility. The workshop was conducted twice – each with half of the participants and lasted for about two hours including a short coffee break. In order to get all participants in the mood for the workshop and to allow them to reflect on the topic from their personal point of view, we invited everybody to complete a brief (~10 min.) "pre-questionnaire" before the workshop (Link: https://thimib.fra1.qualtrics.com/jfe/form/SV_03eUqLNatcDgzs2). For details, see section 3.3. The results from both the questionnaire and the two workshops are currently analyzed and will be later submitted as conference paper or journal article (with recognition of the Dagstuhl Seminar).
- Wednesday, June 8: On this day, in the Dagstuhl tradition to offer a social activity, we watched – again in two groups of each ca. 15 people – the documentary "Coded Bias" (https://www.codedbias.com/). While watching the video, participants were asked to record their thoughts (issues, concerns, surprises, technical problems/solutions, societal/policy related solutions) in a Miro-board, e.g., https://miro.com/app/board/o9J_lCMqEEY=/ for group 1. After watching, we used 10 minutes for clustering the items followed by another 5 minutes for voting. The top voted items where than discussed in the large group and conclusions drawn for our work.

Coded Bias – group 1 results:

- 5 votes: "ensure the right to be forgotten" (removal/deletion of data)
- 4 votes: "AI algorithm uses historical information for the prediction – not everything has been seen before…"
- 3 votes: "Salery automatically based on office environment (stationary, in the car, on the go) -> lot of discussion
- 2 votes: "Transparency"
- 2 votes: "Use a diverse data set to train the AI"
- 2 votes: "Ways of opening the black box…?"

Coded Bias – group 2 results:

- 9 votes: "Transparency and explainability of algorithms (related to and used in automated cars)"
- 6 votes: "Where would bias be exhibited toward passengers or those outside the vehicle?"
- 6 votes: "Lack of regulation and legal structure for AI implementation"
- 6 votes: "mass surveillance unlocked by networked AVs"
- 6 votes: "Ethics education"

- Thursday, June 9: On the second last day of the seminar, all seminar participants met in small groups (2 to 3 people, see Fig. 6.9) to discuss one of the topics identified as most important (and to make a video of the discussion) or to jointly create a Youtube playlist of most-impactful videos in a dedicated topical area related to the seminar. The results were collected by the co-organizers of the seminar and distributed among the participants. Examples of bilateral interviews can be found in Sections 4.1 or 4.2, among others, and an example of a playlist is shown in Section 4.3.
- Friday, June 10: The last day of the seminar has ended with a summit (Fig. 6.10). The first half of this activity was devoted to two panels with distinguished panelists. Panelists started the conversations with brief statements, which were then followed by moderated discussions with the group. For the second half of this activity all participants were sent into breakout rooms in Zoom and worked in smaller groups on a Miro-board (https://miro.com/app/board/o9J_l_-usxU=) on problems discussed during the panels. After the group work, all met again in the main Zoom room and each group presented the results of the group activity (Fig. 6.11).

Fig. 6.4
Group activities on day 1: Collecting of ideas, thoughts, questions from the individual presentations; Majority voting after clustering of collected items.

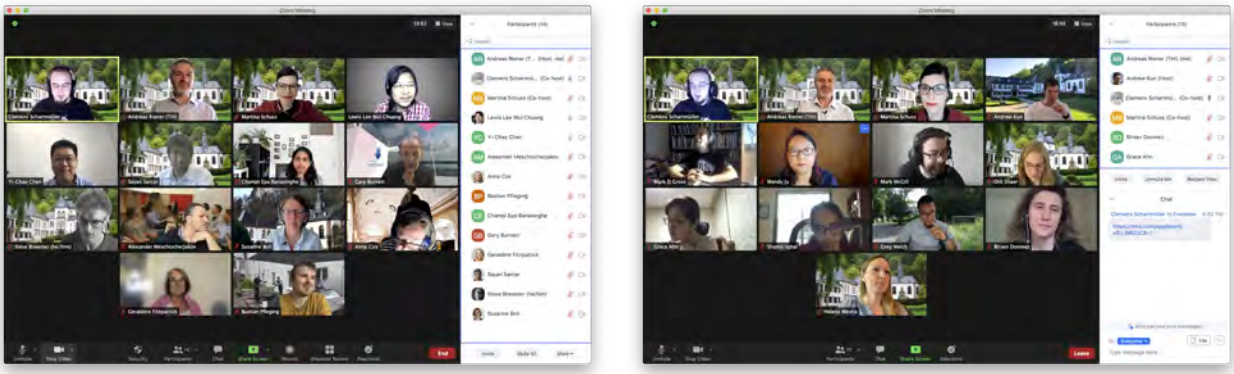

Fig. 6.5
Introduction to the two workshops on day 2 including participants.

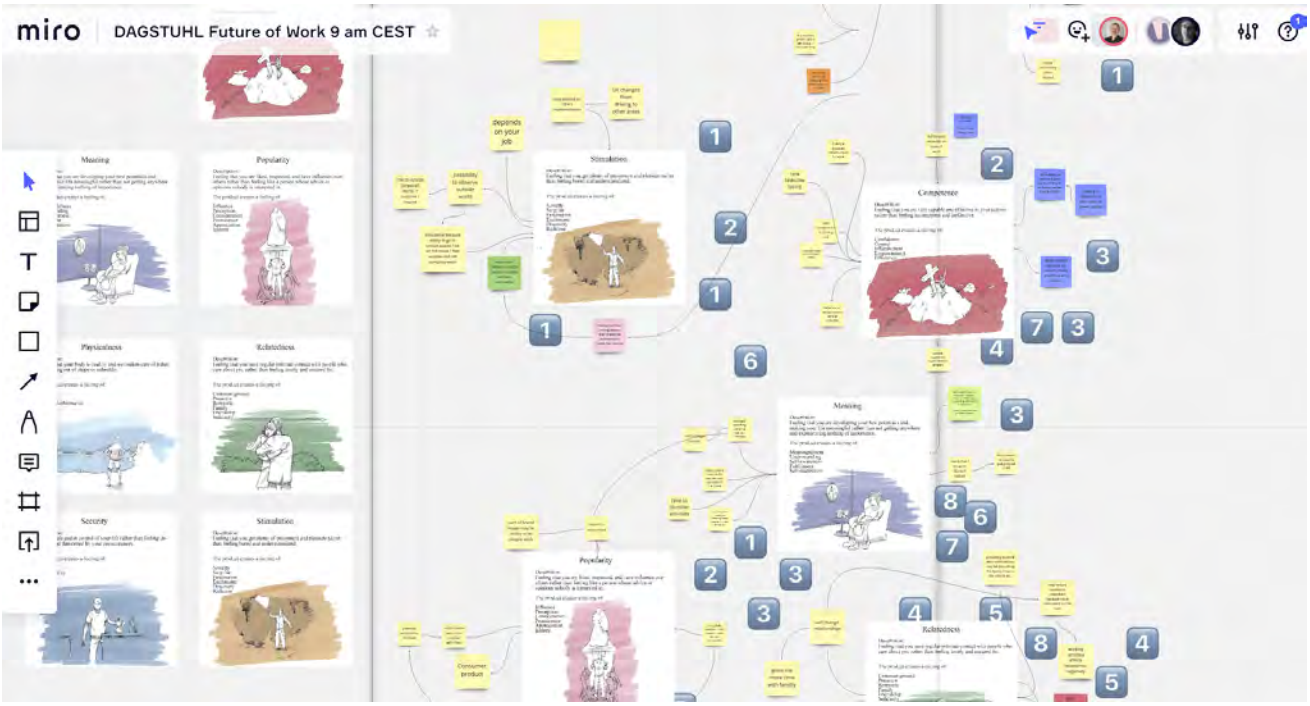

Fig. 6.6
Intermediary results of the interactive workshop part on the Miro-board (group 2 workshop).

Fig. 6.7
Post-its collected by the participants of group 1 and voting results.

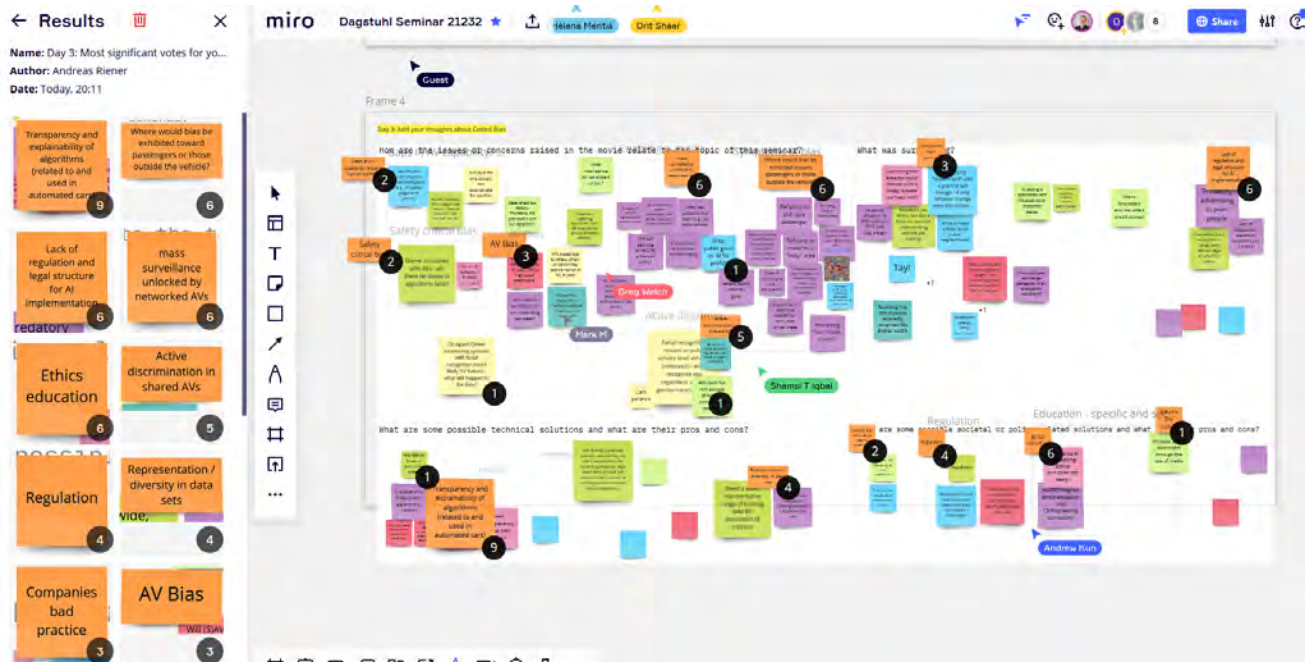

Fig. 6.8
Post-its collected by the participants of group 2 and results of the voting on most relevant elements identified during watching “Coded Bias”.

Dagstuhl Seminar 21232 "Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments"

Participant 1	Affiliation	Participant 2	Affiliation
Chen, Yi-Chao	Shanghai Jiao Tong University	Fitzpatrick, Geraldine	TU Wien
Riener, Andreas	TH Ingolstadt	Burnett, Gary	University of Nottingham
Schartmüller, Clemens	TH Ingolstadt	Ranasinghe, Champika Manel Epa	University of Twente – Enschede
Shaer, Orit	Wellesley College	Lee, John D.	University of Wisconsin – Madison
Boffi, Laura	University of Ferrara	Brewster, Stephen	University of Glasgow
Ahn, Sun Joo	University of Georgia – Athens	Welch, Gregory F.	University of Central Florida – Orlando
Alvarez, Ignacio J.	Intel – Hillsboro	Gross, Mark D.	University of Colorado – Boulder
Boll, Susanne	Universität Oldenburg	McGill, Mark	University of Glasgow
Brumby, Duncan	University College London	Pfleging, Bastian	TU Eindhoven
Meschtscherjakov, Alexander	Universität Salzburg	Cox, Anna	University College London
Chuang, Lewis	IfADo – Dortmund	Lindley, Siân	Microsoft Research – Cambridge
Mentis, Helena M.	University of Maryland – Baltimore County	Ju, Wendy	Cornell Tech – New York
Kun, Andrew	University of New Hampshire – Durham		Sayan Sarkar, Tsukuba University
Donmez, Birsen	University of Toronto	Iqbal, Shamsi Tamara	Microsoft Research – Redmond

Fig. 6.9
Couples who either had a curated conversation or created a Youtube playlist on Thursday bilaterally (<= 5 minutes each).

Fig. 6.10
The highlight of the seminar: A summit with contributions from seminar participants and keynote speeches from invited experts (including Neha Kumar, ACM SIGCHI President).

Fig. 6.11
Overview of the results of the six groups in Miro.

6.11 Compute-First Networking

Organizers: Jon Crowcroft, Philip Eardley, Dirk Kutscher, and Eve M. Schooler

Seminar No. 21243

Date: June 13–16, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.5.54

Remote Participants: Chris Adeniyi-Jones, Laura Al Wardani, Uthra Ambalavanan, Gianni Antichi, Roberto Bifulco, Olivier Bonaventure, Kenneth L. Calvert, Jon Crowcroft, Philip Eardley, T M Rayhan Gias, Tim Harris, Jianfei He, Micchio Honda, Teemu Kärkkäinen, Jussi Kangasharju, Namseok Ko, Michal Król, Ike Kunze, Dirk Kutscher, Julie McCann, Jag Minhas, Marie-Jose Montpetit, Naresh Nayak, Erik Nordmark, David Oran, Jörg Ott, Andy Reid, Eve M. Schooler, Peer Stritzinger, Christian Tschudin, Klaus Wehrle, Cedric Westphal, Peter Willis, Chenren Xu, Noa Zilberman

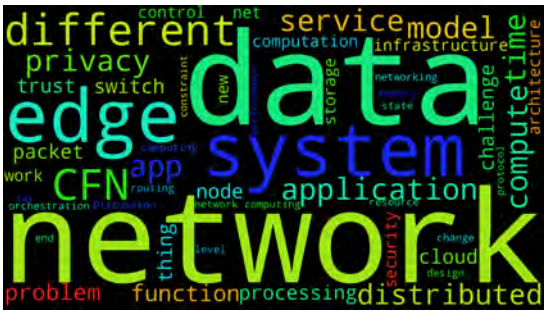

Edge- and more generally In-Network Computing are key elements in many traditional content distribution services today, typically connecting cloud-based computing to consumers. The advent of new programmable hardware platforms, research and wide deployment of distributed computing technologies for data processing, as well as new exciting use cases such as distributed Machine Learning and Metaverse-style ubiquitous computing are now inspiring research of more fine-granular and more principled approaches to distributed computing in the "Edge-To-Cloud Continuum".

The Compute-First Networking Dagstuhl Seminar has brought together researchers and practitioners in the fields of distributed computing, network programmability, Internet of Things, and data analytics to explore the potential, possible technological components, as well as open research questions in an exciting new field that will likely induce a paradigm shift for networking and its relationship with computing.

Traditional overlay-based in-network computing is typically limited to quite specific purposes, for example CDN-style edge computing. At the same time, network programmability approaches such as Software-Defined Networking and corresponding languages such as P4 are often perceived as too limited for application-level programming. Compute-First Networking (CFN) views networking and computing holistically and aims at leveraging network programmability, server- and serverless in-network computing and modern distributed computing abstraction to develop a new system's approach for an environment where computing is not merely and add-on to existing networks, but where networking is re-imagined with a broader and ubiquitous notion of programmability.

We expect this approach to enable several benefits: it can help to unlock distributed computing from the existing silos of individual cloud and CDN platforms – a necessary condition to enable Keiichi Matsuda's vision of Hyper-Reality and Metaverse concepts where the physical world, human users and different forms of analytics, and visual rendering services constantly engage in information exchanges, directly at the edges of the network. It can also help to provide reliable, scalable, privacy-preserving and universally available platforms for Distributed Machine Learning applications that will play a key role in future large-scale data collection and analytics.

CFN's integrated approach allows for several optimizations, for example a more informed and more adaptive resource optimization that can take into account dynamically changing network conditions, availability of utilization of compute platforms as well as application requirements and adaptation boundaries, thus enabling more responsive and better-performing applications.

Several interesting research challenges have been identified that should be addressed in order to realize the CFN vision: How should the different levels of programmability in todays system be integrated into a consistent approach? How would programming and communication abstractions look like? How do orchestration systems need to evolve in order to be usable in these potentially large scale scenarios? How can be guarantee security and privacy properties of a distributed computing slice without having to rely on just location attributes? How would the special requirements and properties of relevant applications such as Distributed Machine best be mapped to CFN – or should distributed data processing for federated or split Machine Learning play a more prominent role in designing CFN abstractions?

This seminar was an important first step in identifying the potential and a first set of interesting new research challenges for re-imaging distributed computing through CFN – an exciting new topic for networking and distributed computing research.

6.12 Quantum Complexity: Theory and Application

Organizers: Bill Fefferman, Sevag Gharibian, Norbert Schuch, and Barbara Terhal

Seminar No. 21261

Date: June 27 to July 2, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.5.76

Participants: Simon Apers, Sergio Boixo, Harry Buhrman, Libor Caha, Jens Eisert, Lior Eldar, Sevag Gharibian, Dominik Hangleiter, Marcel Hinsche, Marios Ioannou, Robert König, Maris Ozols, Dorian Rudolph, Norbert Schuch, Barbara Terhal, Frank Verstraete, Petra Wolf

Remote Participants: Scott Aaronson, Dorit Aharonov, Anurag Anshu, Itai Arad, Johannes Bausch, Adam Bouland, Sergey Bravyi, Michael Bremner, Andrew Childs, Toby Cubitt, Abhinav Deshpande, Bill Fefferman, András Gilyén, David Gosset, Alex Grilo, Sandy Irani, Stacey Jeffery, Elham Kashefi, Joel David Klassen, Robert König, Robin Kothari, Richard Küng, Urmila Mahadev, Milad Marvian, Rewad Mezher, Tomoyuki Morimae, Ramis Movassagh, Daniel Nagaj, Harumichi Nishimura, David Pérez García, Stephen Piddock, Jamie Sikora, Maarten Stroeks, Aarthi Sundaram, Yuki Takeuchi, Ewin Tang, Thomas Vidick, James Watson, Justin Yirka

Background and motivation. Since the seminal discovery of an efficient quantum integer factorization algorithm by Peter Shor in 1994, the field of Quantum Computation has blossomed into a large-scale international effort to build, test, and study the possibilities that information processing using quantum particles may provide. A central role in these developments has been played by Quantum Complexity Theory, a traditionally theoretical realm of research focusing on such questions as: Which physical properties of Nature can be efficiently computed? Can the behavior of an untrusted or noisy quantum computer be verified? What might constitute convincing evidence of "quantum supremacy" over classical computers?

With the first generation of completed "Noisy Intermediate Scale Quantum (NISQ)" experiments already staking quantum supremacy claims, however, the answers to such "traditionally theoretical" questions have taken on an urgent and practical relevance. For example, a complexity theoretic understanding of which realistic physical problems are "just hard enough" for classical computers and "easy enough" for quantum computers is the natural starting point for "quantum supremacy" testbeds. With functioning experimental devices in place, one must next convincingly confirm the device is performing as designed, particularly in the presence of noise. Finally, if the aim of such experiments is to cast doubt on the Extended Church-Turing Thesis, then a strong standard of evidence is required; such a standard must be rigorously stated and developed.

Seminar Topics. This seminar covered a range of topics under the broad umbrella of Quantum Complexity Theory, ranging from highly theoretical to experimentally driven. We briefly overview some of these here; further examples and details are in the included talk abstracts.

Theoretical directions. The field of Quantum Complexity Theory is concerned, broadly speaking, with a rigorous mathematical study of the resources required to perform certain computational tasks. To first order, this involves dividing the "computational world" into two buckets: Easy versus hard problems. However, in reality, the complexity landscape is much finer than this. For example, one might ask – given that problem X has a known efficient quantum algorithm, does there nevertheless exist a *faster* quantum algorithm for X? This typically falls under the classical area of "fine-grained complexity", which has only recently begun to emerge as having a quantum analogue. Conversely, one may ask – is problem X hard only when one wishes to have a

high precision answer, becoming easy when a larger margin of error is allowed? Classically, this falls under the umbrella of "hardness of approximation", and which has seen intense study in the guise of the "quantum PCP conjecture". Finally, given that quantum computers are believed more powerful than classical ones, a natural question is: *Do there exist computational problems whose difficulty lies strictly between classical and quantum?* Here, a natural object of study has been so-called "stoquastic" quantum systems, whose time evolution can often be simulated in practice via randomized (i.e. Monte Carlo) techniques, but which nevertheless appear difficult to classically simulate in the *worst case* in a rigorous fashion. Recent advances and the state of the art in all of these topics, as well as a number of others, were discussed at the seminar.

Experimentally motivated directions. The recent explosion of the so-called Noisy Intermediate-Scale Quantum (NISQ) computation era has brought many new questions to the forefront of Quantum Complexity Theory. For example, to date, two of the leading frameworks for experimental demonstration of "quantum supremacy" have been *random circuit sampling* and *Boson sampling*. On the one hand, much progress has been made closing the remaining gaps in the theoretical hardness proofs for these tasks on classical computers. On the other hand, for experiments that *have* been conducted, important practical topics such as how to benchmark such experimental random circuit setups have very recently been studied. Moreover, beyond the quest for quantum supremacy lies the next question: *What practical applications might NISQ devices already prove useful for?* These and related topics were presented and discussed at the seminar.

Participants and program overview. Due to the on-going COVID situation, the seminar was held in hybrid format. This meant that of the 55 total participants joining from 14 countries around the world (from North America to Europe to Asia), 17 were on site, and 38 were remote. To allow all audience members to participate, a few measures were taken, which arguably worked quite well given the circumstances:

- During each of the seminar's on-site sessions, a Zoom session was projected onto a whiteboard, to which all remote participants were invited. The Zoom participants could see and hear on-site whiteboard and slide presentations, as well as interrupt to ask questions (via the room's loudspeaker system). This made for a reasonably efficient setup in which both on-site and hybrid participants could discuss in real-time. A Slack channel was also set up to ease communication, and by popular request, after talks a virtual Zoom chat room was set up so that the remote participants could also chat amongst themselves.
- To accommodate both types of audience members, a mix of on-site and remote talks were scheduled. On-site talks were typically held in the morning (CEST), allowing remote audience members in Europe Asia to attend. These were held at "standard" times, starting at 9:00 CEST. Remote talks were largely scheduled in the late afternoon and evening (17:00 and 20:00 CEST), this time accessible to North American and European participants.
- Seminar participants Marcel Hinsche (on site) and James Watson (off-site) graciously offered to act as "technical help volunteers" for local and remote participants, ensuring the hybrid setup ran smoothly for both local and remote attendees.

Regarding the remaining program structure, a strong emphasis was placed on plentiful open time for ad-hoc discussion – typically 14:00 to 17:00 was left open expressly for this purpose. A social outing (hike) was organized by participant Dominik Hangleiter on Wednesday afternoon, and a traditional social night in the music room took place on Wednesday evening.

Acknowledgements. The seminar's participants and organizing committee wholeheartedly thank the Schloss Dagstuhl administrative and technical staff, who before, during, and after the seminar were incredibly supportive, professional, and patient with us quantum computer scientists. Many of the seminars participants, both online and off-line, commented very positively of the experience, citing it as a very welcome break to the stress of the on-going COVID pandemic.

6.13 Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users

Organizers: Ana Aguiar, Onur Altintas, Falko Dressler, and Gunnar Karlsson

Seminar No. 21262

Date: June 28, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.5.89

Remote Participants: Ana Aguiar, Onur Altintas, Mate Boban, Claudio Casetti, Wai Chen, Carla-Fabiana Chiasserini, Sinem Coleri, Pedro D'Orey, Klaus David, Falko Dressler, David Eckhoff, Gerhard Fettweis, Julian Heinovski, Frank Hofmann, Gunnar Karlsson, Florian Klingler, Renato Lo Cigno, Ian Robin Marsh, Susana Sargento, Michele Segata, Gürkan Solmaz, Christoph Sommer, Peter Steenkiste, Hsin-Mu Tsai, Elisabeth Uhlemann, Lars Wolf

Looking back at the last decade, one can observe enormous progress in the domain of vehicular networking. In this growing community, many ongoing activities focus on the design of communication protocols to support safety applications, intelligent navigation, cooperative driving and others. Using the terms Vehicular Ad-hoc Networks (VANETs), Inter-Vehicle Communication (IVC), Car-2-X (C2X), or Vehicle-2-X (V2X), many applications – as interesting as challenging – have been envisioned and (at least) partially realized. Very large projects have been initiated to validate the theoretic work in field tests and protocols are being standardized. With the increasing interest from industry, security and privacy have also become crucial aspects in the stage of protocol design in order to support a smooth and carefully planned roll-out. We are now entering an era that might change the game in road traffic management. Many car makers already supply their recent brands with cellular and WiFi modems, some also adding vehicular WLAN (DSRC, ITS-G5) and/or C-V2X technologies, which focus on V2V and V2I communication.

The management and control of network connections among vehicles and between vehicles and an existing network infrastructure is currently one of the most challenging and active research fields in the networking domain. There is a long list of desirable applications that can be grouped into four IVC categories:

1. e-Safety applications that try to make driving safer, e.g., road hazard warning, collision warning;
2. traffic efficiency applications aiming at more efficient and thus greener traffic, e.g., detection of traffic jams, traffic distribution;
3. manufacturer oriented applications, e.g., automatic software updates; and
4. convenience applications, e.g., automatic map updates.

We initiated the "Inter-Vehicular Communication" Dagstuhl Series back in 2010, when a first Dagstuhl Seminar was organized on this topic. The motivation was to bring together experts in this field to investigate the state of the art and to highlight where sufficient solutions already existed. The main outcome of this very inspiring seminar series was that there are indeed areas within this research field where scientific findings are being consolidated and adopted by industry. This was the consensus of quite intriguing discussions among participants from both industry and academia. Yet, even more aspects have been identified where substantial research is still needed.

Some of the findings of the first three seminars in this series have been published not only in the related Dagstuhl reports but also in widely visible magazine articles:

1. Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz and Lars Wischhof, "Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar," IEEE Communications Magazine, vol. 49 (5), pp. 158-164, May 2011.
2. Falko Dressler, Hannes Hartenstein, Onur Altintas and Ozan K. Tonguz, "Inter-Vehicle Communication – Quo Vadis," IEEE Communications Magazine, vol. 52 (6), pp. 170-177, June 2014.
3. Onur Altintas, Suman Banerjee, Falko Dressler and Geert Heijenk, "Executive Summary – Inter-Vehicular Communication Towards Cooperative Driving," Proceedings of Dagstuhl Seminar 18202 on Inter-Vehicular Communication – Towards Cooperative Driving, Schloss Dagstuhl, Germany, May 2018, pp. 31–59.

Seminars in this series focused on general vehicular communication technologies, security and safety impact, cooperative driving concepts and its implications on communication protocol design, and many more.

We now shifted the focus of this seminar from basic networking principles to open challenges in edge computing support and, as a novel aspect, on how to integrate so called vulnerable road users (VRU) into the picture. Edge computing is currently becoming one of the core building blocks of cellular networks, including 5G, and it is necessary to study how to integrate ICT components of moving systems. The trade-offs of computation distribution, system aspects, and the impact on end-to-end latency are still unanswered. Also, vehicular networking and cooperative driving focuses almost exclusively on cars but leaves out communication

and coordination with, for example, pedestrians and bicyclists. For example, many of the existing communication solutions for this scenario were designed without having battery constraints in mind. In the mean-time, some early research has been initiated on this topic, we organized a workshop at INFORMATIK 2019 on VRUs and initial projects report very interesting results on safety features for VRUs. Building upon the great success of the first two seminars, with this follow-up seminar, our goal was to once again bring together experts from all these fields from both academia and industry. The seminar focused intensively on discussions in several working groups. To kick-off these discussions, we invited three keynote talks:

- Distributed machine learning in the vehicle-to-edge continuum by Carla-Fabiana Chiasserini (Polytechnic University of Turin, IT)
- A TechCity Living lab for vehicular-based mobility services in the road by Susana Sargento (Institute of Telecommunications, PT)
- Edge-based increase of awareness and support for all traffic participants by Lars Wolf (TU Braunschweig, DE)

We finally organized the following working groups on some of the most challenging issues related to inter-vehicular communication and cooperative driving:

- Edge Computing: A multi-dimensional techno-economic outlook (i.e., latency, cost, deployment issues, etc.)
- Cooperative Driving Again: Where are we now after 2018? What about vehicular platooning: Is it still alive? What are still unsolved challenges?
- How do we support VRU detection/warning?
- Forget about V2V, we have V2C: Is Vehicle-to-Cloud the Way to Go?

References

1 Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz and Lars Wischhof, “Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar,” IEEE Communications Magazine, vol. 49 (5), pp. 158-164, May 2011.

2 Falko Dressler, Hannes Hartenstein, Onur Altintas and Ozan K. Tonguz, “Inter-Vehicle Communication – Quo Vadis,” IEEE Communications Magazine, vol. 52 (6), pp. 170-177, June 2014.

3 Onur Altintas, Suman Banerjee, Falko Dressler and Geert Heijenk, “Executive Summary – Inter-Vehicular Communication Towards Cooperative Driving,” Proceedings of Dagstuhl Seminar 18202 on Inter-Vehicular Communication – Towards Cooperative Driving, Schloss Dagstuhl, Germany, May 2018, pp. 31–59.

6.14 Computational Proteomics

Organizers: Sebastian Böcker, Rebekah Gundry, Lennart Martens, and Magnus Palmblad
Seminar No. 21271

Date: July 4–9, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.6.1

Participants: Sebastian Böcker, Viktoria Dorfer, Lukas Käll, Frédérique Lisacek, Lennart Martens, Magnus Palmblad, Robin Park, Daniel Questschlich, Veit Schwämmle, Mathias Wilhelm

Remote Participants: Jeffrey Agar, Kiyoko Aoki-Kinoshita, Marshall Bern, Robert Chalkley, Sven Degrove, Bernard Delanghe, Patrick Emery, Rebekah Gundry, Michael Hoopmann, Neil Kelleher, Joanna Kirkpatrick, Daniel Kolarich, Rune Linding, Nicki Packer, Matthew Smith, Sabarinath Peruvemba Subramanian, Morten Thaysen-Andersen, Lilla Turiák, Olga Vitek, Christine Vogel

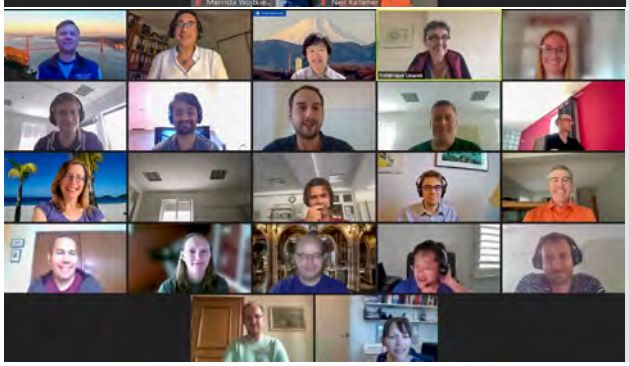

The Dagstuhl Seminar 21271 "Computational Proteomics" discussed several important developments, challenges, and opportunities that are emerging in the field of computational proteomics. Three core topics were set out at the start, and these were discussed at length throughout the seminar.

These three topics were: (i) the fast evolving use of advanced machine learning approaches in proteomics; (ii) the challenges and opportunities offered by fast developing approaches for structural and top-down proteomics; and (iii) specific issues and computational complications in glycoproteomics.

The machine learning and glycoproteomics topics were each introduced by a dedicated lecture, which set out the current state-of-the-art and presented a tentative set of issues, challenges, or opportunities that could be explored during the seminar. The structural and top-down proteomics topic was introduced by two sequential lectures, one on structural proteomics, and one on top-down proteomics. In total, four introductory talks were thus presented at the start of the seminar. For each of the three main topics, daily Working Group sessions were organised, which took place in the morning and afternoon, with a daily late-night

session scheduled each day to wrap up the day's outcomes. This structure was followed to allow maximum involvement by online participants across the various timezones in the hybrid format. The Machine Learning in Proteomics Working Group also spun out another Working Group session during the seminar, which discussed the creation of a machine learning (Kaggle-like) competition based on proteomics data.

Each of these breakout sessions was very actively attended, including by online attendees, and resulted in several interesting research ideas and potential new initiatives. The Machine Learning in Proteomics Working Group was the largest working group, and addressed a number of distinct topics during the seminar. Of particular note were the spin-out effort to establish two machine learning competitions based on proteomics data and challenges to engage the broader machine learning community, and the extensive discussions on the optimal way to represent mass spectrometry data for downstream machine learning.

The Glycoproteomics Working Group was very actively attended, and discussed an exciting set of topics. A first highlight among these topics was provided by the extensive and detailed discussions with the Machine Learning Working Group regarding the potential of, and road towards, the use of state-of-the-art machine learning approaches in glycoproteomics. A second highlight concerned the delineation of a set of high-impact opinion papers to describe the state-of-the-art of the field, and its goals, ambitions, and challenges.

The Structural and Top-Down Proteomics Working Group was very active in detailing the many challenges and opportunities in this fast-evolving field. One noteworthy challenge revolved around the detection, annotation, and biological interpretation of post-translational modifications detected by mass spectrometry. A second challenge concerned the standardization of acquired native mass spectrometry data, the minimal reporting requirements for these experiments, and the dissemination of these data.

Overall, the 2021 Dagstuhl Seminar on Computational Proteomics was extremely successful as a catalyst for careful yet original thinking about key challenges in the field, and as a means to enable downstream progress by setting important, high impact goals to work on in close collaboration. During this Seminar, new topics for a future Seminar were suggested throughout as well, indicating that this active field will continue to yield novel challenges and opportunities for advanced computational work going forward.

6.15 Towards Climate-Friendly Internet Research

Organizers: Vaibhav Bajpai, Jon Crowcroft, Oliver Hohlfeld, and Srinivasan Keshav

Seminar No. 21272

Date: July 6–9, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.6.14

Remote Participants: Jari Arkko, Vaibhav Bajpai, Sujata Banerjee, Georg Carle, Jon Crowcroft, Quentin De Coninck, Simone Ferlin, Andrew Hines, Oliver Hohlfeld, Daniel Karrenberg, Wolfgang Kellerer, Srinivasan Keshav, Mirja Kühlewind, Mirjam Kühne, Franziska Lichtblau, Michael Menth, Jörg Ott, Cristel Pelsser, Colin Perkins, Alexander Raake, Amr Rizk, Jürgen Schönwälder, Henning Schulzrinne, Georgios Smaragdakis, Ralf Steinmetz, Cristina Videira Lopes, Martina Zitterbart

Goals

The Internet was originally developed to ease collaboration between remote parties, thereby, in principle, reducing carbon emissions by a reduced need for travel. Yet, conducting research on communication networks has typically involved a certain level of carbon footprint. One fundamental reason is the publication and dissemination culture in the field, which focuses on conferences and workshops rather than journals. Not only does every dissemination of a research result therefore involves travel, even the peer-review process to decide which papers to accept, in the form of an in-person technical program committee (TPC) meeting, also requires travel. Moreover, although the standardization of Internet technology within the Internet Engineering Task Force (IETF) largely involves online discussions and audio/video streaming–unlike almost all other standardization bodies–yet regular in-person meetings are considered critical to converge discussion and build consensus. Thus, conducting and disseminating networking research has resulted in a high level of travel, and a consequent high carbon footprint.

The carbon footprint of these trips (mostly air travel) can, however, be reduced by means of organizational changes and virtual conferences. Recently, as a consequence of the COVID-19 pandemic, we have already witnessed a rapid transition to a virtual mode of operation including remote working, online meetings, and virtual conferences. This has resulted in first-hand experience in carrying out research but with no travel.

In this Dagstuhl Seminar, we initiated a discussion on how to make Internet research more climate friendly. Specifically, we evaluated experiences in running and participating in virtual conferences as a consequence of the COVID-19 pandemic. We wanted to understand what went well and what went badly in implementing and deploying virtual conferences, what challenges were encountered, and what needs to be improved, particularly as we transition to hybrid in-person, online meetings. The broader goal of the seminar is to identify how to transition to a new status quo that continues to reduce the carbon footprint from travel.

Structure

The seminar lasted 2.5 days. It began with an introductory round where each participant presented one slide to give an overview of their experience that was relevant for the seminar and a set of open questions that the participant wished to discuss during the event. These slides were collected from each participant before the seminar. We also had one invited talk (§3.1 of the full report) that we used as a basis for triggering discussions and identifying areas for group work, while a major portion of the seminar time was dedicated to breakout sessions, whereby participants were split into small groups to discuss specific themes and develop ideas with consensus to propose to larger groups. The morning sessions the following day were dedicated to continuing parallel group work with presentations that reported the outcomes of each breakout session from the previous day. Every evening, we had an online social activity. The afternoon of the third day was spent reviewing and collecting feedback from the participants and for initiating follow up actions identified during the seminar.

6.16 Data Structures for Modern Memory and Storage Hierarchies

Organizers: Stratos Idreos, Viktor Leis, Kai-Uwe Sattler, and Margo Seltzer

Seminar No. 21283

Date: July 11–16, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.6.38

Participants: Gustavo Alonso, Alexander Baumstark, Carsten Binnig, André Brinkmann, Christian Dietrich, Muhammad Attahir Jibril, Alfons Kemper, Viktor Leis, Alberto Lerner, Ulrich Carsten Meyer, Thomas Neumann, Ismail Oukid, Marcus Paradies, Kai-Uwe Sattler, Jens Teubner, Alexander van Renen

Remote Participants: Marcos K. Aguilera, Raja Appuswamy, Manos Athanassoulis, Alexander Böhm, Peter A. Boncz, Mark Callaghan, Khuzaima Daudjee, Jana Giceva, Goetz Graefe, Gabriel Haas, Stratos Idreos, Wolfgang Lehner, Danica Porobic, Ken Salem, Wolfgang Schröder-Preikschat, Margo Seltzer, Tianzheng Wang, William Wang

The seminar brought together researchers and practitioners from the data management and systems/storage communities to discuss the implications of the modern hardware landscape on high-performance systems. Due to the pandemic, the seminar was organized as a hybrid event: Virtual participation was limited to one session per day that featured invited talks. The in-person component consisted of free-flowing plenary discussions and several smaller, focused working groups. Some key takeaways from the discussion are:

- **OS/DBMS co-design**: Traditional POSIX-style OS abstractions do not work well for data-intensive systems, leading to complex workarounds and suboptimal performance. While some of these issues could in principle be fixed by optimizing OS implementations, others require new APIs. For example, it is very difficult to implement crash-consistent data structures on top of the mmap system call.
- **Cloud**: The cloud is taking over and cloud-native data processing systems often have a a very different architecture from traditional data management systems. For example, many systems strive to separate storage from compute. This trend is enabled by ever faster networks.
- **Near-data processing**: Separating storage from compute leads to costly data movement, which may be mitigated by pushing down (parts of) the computation close to the data. Major public cloud vendors already to optimize their internal services towards this goal. The challenges is how to program such distributed and specialized hardware components.
- **Persistent Memory**: One major question discussed at the seminar was the role of byte-addressable persistent memory in future systems and whether what the "kill app" for this technology is. While there are several promising applications (e.g., graph processing or systems that require fast recovery times), it is not clear whether wide adoption will occur. Currently, the technology is quite expensive (prices per byte are similar to DRAM) and very hard to program in a crash-consistent way (e.g., writes must be carefully ordered similar to lock-free-style programming).

6.17 Scalable Handling of Effects

Organizers: Danel Ahman, Amal Ahmed, Sam Lindley, and Andreas Rossberg

Seminar No. 21292

Date: July 18–23, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.7.54

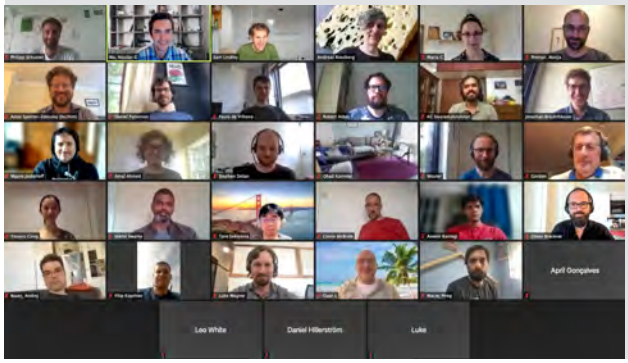

Remote Participants: Danel Ahman, Amal Ahmed, Robert Atkey, Andrej Bauer, Oliver Bracevac, Jonathan Immanuel Brachthäuser, Youyou Cong, Paulo Emílio de Vilhena, Stephen Dolan, Ronald Garcia, April Gonçalves, Maria Gorinova, Daniel Hillerström, Mauro Jaskelioff, Ohad Kammar, Oleg Kiselyov, Filip Koprivec, Sivaramakrishnan Krishnamoorthy Chandrasekaran Indian Institute of Techology, IN, Daan Leijen, Sam Lindley, Conor McBride, Daniel Patterson, Maciej Piróg, Gordon Plotkin, Matija Pretnar, Aseem Rastogi, Andreas Rossberg, Philipp Schuster, Taro Sekiyama, Antal Spector-Zabusky, Nikhil Swamy, Wouter Swierstra, Luke Wagner, Leo White, Nicolas Wu

Algebraic effects and effect handlers are currently enjoying significant interest in both academia and industry as a modular programming abstraction for expressing and incorporating user-defined computational effects in programming languages. For example, there are a number of effect handler oriented languages in development (such as Eff, Frank, and Koka); there exist effect handler libraries for mainstream languages (such as C and Java); effect handlers are a key part of languages such as Multicore OCaml (and indeed they are due to appear in the production release of OCaml next year); effect handlers are being increasingly used in statistical probabilistic programming (such as Uber's Pyro tool); and proposals are in the works to include effect handlers in new low-level languages (such as WebAssembly). While effect handlers have solid mathematical foundations and have been extensively experimented in prototype languages and on smaller examples, enabling effect handlers to scale still requires tackling some hard problems. To this end, this Dagstuhl Seminar 21292 "Scalable Handling of Effects" focused on addressing the following key problem areas for scalability: Safety, Modularity, Interoperability, Legibility, and Efficiency.

This seminar followed the earlier successful Dagstuhl Seminars 16112 "From Theory to Practice of Algebraic Effects and Handlers" and 18172 "Algebraic Effect Handlers go Mainstream", which were respectively dedicated to the foundations of algebraic effects and to the introduction of them into mainstream languages. In contrast to these previous two seminars which took place in person at Schloss Dagstuhl, the current seminar was organised fully online due to the SARS-CoV-2 pandemic. As the seminar was attended by participants from a wide range of time zones (ranging from the West coast of the US all the way to Japan), coming up with a schedule that was suitable for everybody was a challenge. In the end, we decided to have three scheduled two-hour sessions each day, with impromptu informal discussions also happening between-times. These sessions were: (i) 15:00-17:00 CEST, which were deemed the Core Hours, where all participants were most likely to be able to present; (ii) 10:00-12:00 CEST, which was most suitable for participants from Asia and Europe; and (iii) 17:30-19:30 CEST, which was most suitable for participants from America and Europe. The Core Hours included talks, breakouts, and discussions of interest to the widest audience, with more specialised talks and breakouts taking place in the other two daily scheduled blocks. Talks were recorded so that participants could catch up due to being in an incompatible time zone, then deleted at the end of the week.

In order to run a successful virtual Dagstuhl Seminar we exploited several different technologies. For talks we used Zoom. For breakouts we used a combination of Zoom and Gather.town, and for asynchronous communication and further discussions we used Zulip. For scheduling purposes, we used the wiki page provided by Dagstuhl.

We collected initial lists of proposed talks and breakout topics before the seminar began using an online form. We extended these throughout the week. We scheduled talks and breakout groups daily depending on audience interest and the participant availability. While the first part of the week was dominated by talks, the second part of the week saw more emphasis on breakouts and discussions. During Friday's Core Hours, the leaders of each breakout group presented a short overview of the discussions and results (11 reports in total). Initially, we were a little unsure about how well breakout sessions would work in a virtual seminar, but as the week went on they became more and more popular and they seemed to go remarkably well. Initially, we mostly used Gather.town and its virtual whiteboards for the breakout sessions. Subsequently, we transitioned to mostly using Zoom breakout rooms (partly because some people had difficulty using Gather.town on their systems).

The seminar was a great success, particularly given the constraints of the virtual format.

There were vibrant discussions around multishot continuations. These are vital for exciting new applications such as probabilistic programming and automatic differentiation, but more research is needed on how to implement them safely and efficiently in different contexts. Flipping perspective, it was mooted that for certain applications, particularly those involving

direct interaction with the external world, it might be worthwhile restricting attention to runners, which are even more constrained than effect handlers with singleshot continuations.

There were several discussions relating to usability of effect handlers. These resulted in proposals to design a lecture course on effect handlers and to write a book on how to design effectful programs.

A major area of interest instigated at a prior Dagstuhl Seminar (18172 "Algebraic Effect Handlers go Mainstream") is the addition of effect handlers to WebAssembly. A design is being actively worked on as part of the official WebAssembly development process. At the current seminar we worked out extensions to the existing proposal to accommodate named effect handlers and symmetric stack-switching, both of which promise more efficient execution.

An issue with many existing benchmarks for effect handlers is that they often require installing a range of experimental software and configuring it with just the right settings. In order to make it easier to compare systems and share experimental setups we created the effect handlers benchmarks suite – a repository of benchmarks and systems covering effects and handlers in various programming languages, based on Docker scripts that make it easy for anyone to run the benchmarks and adapt them for their own research. The repository is hosted on GitHub. Since the seminar, 5 systems have been added to the repository and it has been actively updated and maintained by different members of the community.

At the end of the week, there was strong interest among the participants to continue this successful seminar series and submit a proposal for another incarnation, hopefully possible to take place on site in about two years.

6.18 Parameterized Complexity in Graph Drawing

Organizers: Robert Ganian, Fabrizio Montecchiani, Martin Nöllenburg, and Meirav Zehavi

Seminar No. 21293

Date: July 18–23, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.6.82

Participants: Michael A. Bekos, Steven Chaplick, Giordano Da Lozzo, Emilio Di Giacomo, Walter Didimo, Fabrizio Frati, Robert Ganian, Martin Gronemann, Thekla Hamm, Petr Hlinený, Michael Kaufmann, Philipp Kindermann, Boris Klemz, Stephen G. Kobourov, Giuseppe Liotta, Maarten Löffler, Fabrizio Montecchiani, Martin Nöllenburg, Chrysanthi Raftopoulou, Ignaz Rutter, Kirill Simonov, Manuel Sorge, Birgit Vogtenhuber, Alexander Wolff, Jules Wulms, Johannes Zink

Remote Participants: Meirav Zehavi, Siddharth Gupta

Graph Drawing. Graph-based models are pervasive in many fields of science and technology. Very often scientists and users analyze these models and communicate their findings by means of graphical representations. This motivated the birth and evolution of *graph drawing*, a self-standing discipline that has evolved tremendously over the past 50 years. Today graph drawing is a mature area of computer science [5, 13, 17, 18] with its own annual conference, the International Symposium on Graph Drawing and Network Visualization (GD)[41]. The focus of the research area today is on combinatorial and algorithmic aspects of drawing graphs as well as on the design of network visualization systems and interfaces. Graph drawing is motivated by applications where it is crucial to visually analyze and interact with relational datasets. Examples of such application areas include data science, social sciences, web computing, information systems, biology, geography, business intelligence, information security and software engineering.

Roughly speaking, graph drawing deals with the construction and analysis of geometric representations of graphs and networks subject to specific layout conventions, such as different notions of planarity or more general crossing constraints, grid layouts, orthogonal drawings etc. Many classic graph drawing problems are NP-hard and thus a variety of theoretical and practical algorithmic techniques for dealing with hard problems are required in graph drawing.

Parameterized Complexity. Numerous computational problems of wide interest are known to be NP-hard in general. Yet, it is often possible to utilize the structure implicitly underlying many real-world instances to find exact solutions efficiently. There is long-standing systematic research of tractability results for various problems on specific classes of instances, and research in this direction constitutes one of the fundamental areas of computer science. However, in many real-world situations it is not possible to define a clear-cut class of instances that we wish to solve; instead of being black and white (belonging to a specific class or not), instances often come in various shades of grey (having certain degrees of internal structure).

The relatively young *parameterized complexity* paradigm [4, 6, 8, 16] offers the perfect tools to deal with this situation. In the parameterized setting, we associate each instance with a numerical *parameter*, which captures how "structured" the instance is. This then allows the development of algorithms whose performance strongly depends on the parameter – instead of the classical setting, where we often associate tractability with polynomial running times and intractability with superpolynomial ones, parameterized algorithms naturally "scale" with the amount of structure contained in the instance. The central notion of tractability in the parameterized setting is *fixed-parameter tractable* (FPT in short), which means that the given problem can be solved by an algorithm with runtime of the form $f(k) \cdot n^{\mathcal{O}(1)}$ (where f is an arbitrary computable function, k is the value of the parameter, and n is the input size). Aside from fixed-parameter tractability, the parameterized complexity landscape consists of a variety of companion notions such as *XP-tractability*, *kernelization* and *W-hardness*.

Parameterized Complexity in Graph Drawing. Research at the intersection of graph drawing and parameterized complexity (and parameterized algorithms in particular) is in its

[41] see http://www.graphdrawing.org

infancy. Most of the early efforts have been directed at variants of the classic Crossing Minimization problem, introduced by Turán in 1940 [19], parameterized by the number of crossings. Here, the objective is to draw a given graph in the plane so as to induce minimum number of crossings. Already in 2001, it was shown to be FPT [9]. A few subsequent works followed [11, 14], including the best paper of GD 2019 [12], but also concerning restricted layouts such as two-layered embeddings [7] and two-sided circular graph layouts [15]. On a related note, given a graph drawn in the plane, some preliminary works considered the detection of a subgraph having a particular structure with minimum number of crossings [1, 10]. Recently, parameterized analysis of specific embeddings such as book embeddings [2, 3], was also brought into life. Overall, the intersection of graph drawing and parameterized complexity still remains mostly unexplored, yet we see many interesting challenges and opportunities for taking a parameterized perspective on graph drawing problems and investigating the applicability of advanced parameterized techniques.

Seminar Goals

The main goal of the seminar was to chart new paths towards research combining the latest findings and techniques in parameterized complexity and graph drawing. In particular, the seminar focused on several prominent topics in graph drawing as well as state-of-the-art tools in parameterized complexity. The discussions addressed both concrete open problems as well as general directions for future research. An integral part of these discussions was the identification and formulation of major challenges as well as novel parameterizations of graph drawing problems relevant to parameterized analysis. The discussions also addressed the applicability of classic as well as cutting-edge tools in parameterized complexity to graph drawing.

In view of the above, it is safe to say that the selection of suitable problems to target was of great importance for the success of the seminar. Our main aim was to offer the participants the opportunity to propose problems to work on, and so the final selection of problems targeted by working groups was carried out during the seminar itself. That being said, we have also prepared a list of candidate problems that we believe would be prime candidates for further investigation through the lens of parameterized complexity.

Seminar Program

1. On the first day of the seminar we enjoyed short introductions of all participants, and four invited overview lectures on different research domains within Graph Drawing. The topics and speakers were chosen as to create a joint understanding of the state of the art of problems in Graph Drawing suitable for parameterized anaylsis. Thekla Hamm presented the topic of graph drawing extension problems, Petr Hliněný presented the topic of planar insertion problems, Michael Kaufmann presented the topic of graph drawing beyond planarity and parametrized complexity, and Ignaz Rutter presented the topic of constrained embedding problems. More information on each lecture can be found in Section 3 of the full report. Overall, this day prepared the ground for the open problem session on the second day.
2. The open problem session took place in the morning of the second day of the seminar. In this session, we collected a list of open research problems that were contributed by the seminar participants. In a preference voting we determined the five topics that raised the most interest among the participants and formed small working groups around them. Each group contained experts in both Graph Drawing and Parameterized Complexity. During the following days the groups worked by themselves, except for a few plenary reporting sessions, formalizing and solving their respective challenges. Below is a list of the working group topics; more detailed group reports are found in Section 5 of the full report.
   a. **Upward/level planarity:** This group studied two previously established restrictions of drawing planar graphs: vertices are either assigned a "horizontal level" that they must be placed on, or there are directed arcs and the drawing must have all edges facing upwards. The group aimed at the development of new parameterized algorithms for both of these NP-hard problems.
   b. **Two-page embeddings of upward planar graphs:** The group studied the complexity of recognizing whether st-planar graphs admit an upward two-page book embedding.
   c. **Orthogonal drawings:** The group focused on the Compaction problem (computing a minimum-area drawing for an orthogonal graph), parameterized primarily by the number of kitty corners, that is, pairs of reflex vertices that point to each other.
   d. **Almost Separated Fixed Order Stack Layouts:** In a fixed order stack layout the vertices of a graph are given with a fixed order and one has to assign the edges to pages so that no two edges on any page cross. This group studied a variant of this well-known NP-complete problem, where the graph is bipartite and the vertices form k consecutive blocks from either part.
   e. **Graph product structure theorem:** The group considered strong products of graphs that yield supergraphs of k-planar graphs, i.e., of graphs that admit a drawing in the plane in which each edge is crossed at most k times. The objective is to exploit these products to derive new upper bounds on the queue number of k-planar graphs.
   f. **Decision trees:** Decision Trees are well known tools used to describe, classify, and generalize data. Besides their simplicity, decision trees are particularly attractive for providing interpretable models of the underlying data. The group studied the complexity of learning decision trees of minimum size under several different parameterizations.
3. After the open problem session, Robert Ganian gave a tutorial in the second day of the seminar that showcased how some of the tools in Parameterized Complexity can be applied to difficult problems, with a special focus on problems that are relevant to graph drawing. The tutorial was prepared in a way so as to make it accessible to the graph drawing community, acting as catalysis for progress on the five selected topics.
4. In the rest of the second day and the other days of the seminar, we had a flexible working schedule with a short plenary session every morning to accommodate group reports and impromptu presentations by participants.

Future Plans

The seminar was designed to foster new research collaborations between researchers in the graph drawing and parameterized complexity communities, whose paths rarely cross in the traditional conferences. These collaborations are very likely to result in new breakthroughs and results, and we expect that the seminar will lead to tangible progress in our understanding of problems of interest. In this sense, the primary outcome from the

seminar will be research papers published at the core conferences and journals for the graph drawing and parameterized complexity communities, such as:

- The International Symposium on Computational Geometry (**SoCG**),
- The International Symposium on Graph Drawing and Network Visualization (**GD**),
- The ACM-SIAM Symposium on Discrete Algorithms (**SODA**), and
- The International Symposium on Theoretical Aspects of Computer Science (**STACS**).

In the mid- and long-term horizon, the seminar will also help build a bridge between the two communities and identify other interesting graph drawing problems which would benefit from a rigorous investigation using tools from parameterized complexity. It can also lead to the development of new parameterized tools and techniques that are designed to deal with the specific obstacles that arise when trying to apply parameterized approaches in the graph drawing setting. Last but not least, the seminar will raise the awareness for the typical research problems and the latest techniques in each others community and thus enrich the knowledge and toolbox of individual participants.

Dagstuhl Seminar in 2022/2023 on Graph Drawing in Parameterized Complexity. This Dagstuhl Seminar has revealed, for the first time in a systematic way, the astounding wealth of problems in Graph Drawing that are naturally multivariate and hence suitable for parameterized analysis; thus a follow-up Dagstuhl Seminar will be proposed to further discuss and deepen our understanding of this topic whose full potential is yet to be unlocked, once again bringing together researchers in Graph Drawing and Parameterized Complexity.

Evaluation

According to the Dagstuhl survey conducted after the seminar, as well as informal feedback to the organizers, the seminar was highly appreciated. Particularly the small group size, group composition, and the seminar structure focusing on hands-on working groups was very well received. The seminar's goals to identify new research directions and initiate collaborations at the intersection of the two different fields of Graph Drawing and Parameterized Complexity was very successful (also in comparison to other Dagstuhl Seminars). Indeed, the participants rated the seminar highly for the mixture of these two fields and its productive interdisciplinary atmosphere, yielding new research perspectives, which have also resulted in new collaborations, joint projects and publications. We are looking forward to seeing the first scientific outcomes of the seminar in the near future and to continuing the efforts to support the growth of interest in parameterized analysis of problems in Graph Drawing.

The seminar had more participants from the Graph Drawing community than from the Parameterized Complexity community due to critical uncertainties caused by the COVID-19 pandemic. We hope that the current trend of improvement in the situation will help in composing a more balanced list of participants in future seminars on this topic.

Acknowledgments

Schloss Dagstuhl was the perfect place for hosting a seminar like this. The unique scientific atmosphere and the historic building provided not only all the room we needed for our program and the working groups, but also plenty of opportunities for continued discussions and socializing outside the official program, especially in these difficult times during the COVID-19 pandemic with all participants being eager to meet and do research together in real life. On behalf of all participants, the organizers want to express their deep gratitude to the entire Dagstuhl staff for their outstanding support and service accompanying this seminar. We further thank Jules Wulms for helping us collect the contributions and prepare the report.

References

1 Akanksha Agrawal, Grzegorz Guspiel, Jayakrishnan Madathil, Saket Saurabh, and Meirav Zehavi. Connecting the Dots (with Minimum Crossings). In Gill Barequet and Yusu Wang, editors, *Symposium on Computational Geometry (SoCG 2019)*, volume 129 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 7:1–7:17, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

2 Michael J. Bannister, David Eppstein, and Joseph A. Simons. Fixed parameter tractability of crossing minimization of almost-trees. In *Graph Drawing (GD 2013)*, volume 8242 of *Lecture Notes in Computer Science*, pages 340–351. Springer, 2013.

3 Sujoy Bhore, Robert Ganian, Fabrizio Montecchiani, and Martin Nöllenburg. Parameterized algorithms for book embedding problems. In *Graph Drawing and Network Visualization (GD 2019)*, Lecture Notes in Computer Science. Springer, 2019. To appear.

4 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.

5 Giuseppe Di Battista, Peter Eades, Roberto Tamassia, and Ioannis G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice-Hall, 1999.

6 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer Verlag, 2013.

7 Vida Dujmovic, Michael R. Fellows, Matthew Kitching, Giuseppe Liotta, Catherine McCartin, Naomi Nishimura, Prabhakar Ragde, Frances A. Rosamond, Sue Whitesides, and David R. Wood. On the parameterized complexity of layered graph drawing. *Algorithmica*, 52(2):267–292, 2008.

8 F.V. Fomin, D. Lokshtanov, S. Saurabh, and M. Zehavi. *Kernelization: Theory of Parameterized Preprocessing*. Cambridge University Press, 2018.

9 Martin Grohe. Computing crossing numbers in quadratic time. *J. Comput. Syst. Sci.*, 68(2):285–302, 2004.

10 Magnús M. Halldórsson, Christian Knauer, Andreas Spillner, and Takeshi Tokuyama. Fixed-parameter tractability for non-crossing spanning trees. In *Algorithms and Data Structures (WADS 2007)*, volume 4619

of *Lecture Notes in Computer Science*, pages 410–421. Springer, 2007.

11 Petr Hlinený and Marek Dernár. Crossing number is hard for kernelization. In *Symposium on Computational Geometry (SoCG 2016)*, volume 51 of *LIPIcs*, pages 42:1–42:10. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2016.

12 Petr Hlinený and Abhisekh Sankaran. Exact crossing number parameterized by vertex cover. In *Graph Drawing and Network Visualization (GD 2019)*, Lecture Notes in Computer Science. Springer, 2019. To appear.

13 Michael Jünger and Petra Mutzel, editors. *Graph Drawing Software*. Springer, 2004.

14 Ken-ichi Kawarabayashi and Bruce A. Reed. Computing crossing number in linear time. In *Symposium on Theory of Computing (STOC 2007)*, pages 382–390. ACM, 2007.

15 Fabian Klute and Martin Nöllenburg. Minimizing crossings in constrained two-sided circular graph layouts. In *Symposium on Computational Geometry (SoCG 2018)*, volume 99 of *LIPIcs*, pages 53:1–53:14. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2018.

16 Rolf Niedermeier. *Invitation to Fixed-Parameter Algorithms*. Oxford Lecture Series in Mathematics and Its Applications. OUP Oxford, 2006.

17 Takao Nishizeki and Md. Saidur Rahman. *Planar Graph Drawing*, volume 12 of *Lecture Notes Series on Computing*. World Scientific, 2004.

18 Roberto Tamassia, editor. *Handbook on Graph Drawing and Visualization*. Chapman and Hall/CRC, 2013.

19 Paul Turán. A note of welcome. *Journal of Graph Theory*, 1(1):7–9, 1977.

6.19 Matching Under Preferences: Theory and Practice

Organizers: Haris Aziz, Péter Biró, Tamás Fleiner, and Bettina Klaus

Seminar No. 21301

Date: July 25–30, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.6.124

Participants: Péter Biró, Somouaoga Bonkoungou, Florian Brandl, Jiehua Chen, Ágnes Cseh, Lars Ehlers, Tamás Fleiner, Martin Hoefer, Zsuzsanna Jankó, Bettina Klaus, Simon Mauras, Seckin Özbilen, Katarzyna Paluch, Jan Christoph Schlegel, Karolina Lena Johanna Vocke

Remote Participants: Haris Aziz, Martin Bichler, Felix Brandt, Robert Bredereck, Christine Cheng, Sanmay Das, John Dickerson, Di Feng, Sushmita Gupta, Klaus Heeger, Hadi Hosseini, Pallavi Jain, Yash Kanoria, Flip Klijn, Fuhito Kojima, Alexander Lam, Irene Y. Lo, Nicholas Mattei, Michael McKay, Shuichi Miyazaki, Thayer Morrill, Thanh Nguyen, Sofiat Olaosebikan, Daniel Paulusma, Baharak Rastegari, Ildikó Schlotter, Jay Sethuraman, Kavitha Telikepalli, Alexander Teytelboym, Utku Unver, Toby Walsh, Mobin YahyazadehJeloudar, Yu Yokoi, Makoto Yokoo

Matching under preferences is a general field spanning computer science, economics, and mathematics. The seminal paper in the field is one by Gale and Shapley (1962) that launched an algorithmic approach to matching agents with preferences. The central problems in the field involve matching agents to each other and to resources in a stable and efficient manner. Matching market algorithms based on the preferences of the agents have several applications such as in school admissions, placement of hospital residents, and centralized kidney markets. Topics in the field include two-sided matchings involving agents on both sides (e.g., job markets, school choice, etc.); two-sided matchings involving agents and items (e.g., course allocation, project allocation, assigning papers to reviewers etc.); one-sided matchings (roommates problem, kidney exchanges, etc.); and matching with payments (assignment game, auctions, etc.).

The topic of matching under preferences not only has tremendous applications but is based on a deep mathematical theory that has been developed by multiple research communities including theoretical computer science, artificial intelligence, discrete mathematics, game theory, and microeconomics. One of the main purposes of the seminar was to bring together leading researchers from various communities working on the topic and facilitate collaboration. The participant list was a mixture of researchers from computer science, mathematics, and economics. The seminar provided a platform to discuss state of the art in matching under preferences; identify new and exciting applications of developing research; and understand the mathematical and algorithmic requirements of new and upcoming problems in the field.

The seminar was conducted in a hybrid manner, with 15 participants attending the seminar physically from the Dagstuhl center and 34 participants attending online. The hybrid nature of the work required the need for careful planning to keep participants engaged and to facilitate collaboration between off-site and on-site participants. The online participation was managed via zoom and gather-town softwares.

The four main focus topics of the seminar were the following ones.

1. Matching markets with distributional constraints,
2. Probabilistic and Fractional Matching,
3. Matching in online and dynamic settings, and
4. Matching Markets and machine learning.

All of the four focus areas are important directions for the field. As new applications arise, it is clear that many real-life matching markets have additional feasibility and distributional constraints. Secondly, most of the theoretical developments in the field concern deterministic outcomes, so one of the goals was

to make further progress on probabilistic mechanisms. During the seminar, current and new research on probabilistic approaches was discussed. Thirdly, many practical matching markets have online and dynamic aspects. There were several discussions on how to model and solve online matching problems. Fourthly, with the increased importance of machine learning in building computer systems, the seminar provided an opportunity to discuss how learning approaches help solve market design problems.

On each of the first four days, there was a one-hour survey talk given by an expert on the above topics. On the first day, Yash Kanoria presented a survey on "online matching markets". On the second day, John Dickerson gave a survey talk on machine learning and matching markets. On day three, Makoto Yokoo presented an overview of "matching under constraints." On day four, Jay Sethuraman surveyed "probabilistic matchings."

On each of the days there were several shorter scientific presentations. During the seminar, two rump sessions were organized to facilitate different time zones. The rump sessions gave an opportunity to the participants to give brief updates or share open problems.

During the week, there were several time slots dedicated to flexible discussion and collaboration as well as dedicated working groups working on particular research topics. Apart from collaborations in smaller groups, the seminar witnessed major collaboration or discussion around several topics. Robert Bredereck brought up the issue of unifying and streamlining terminology and discussed the issue of using gender-neutral terms. There was a large working group led by Sushmita Gupta and Pallavi Jain that examined computational problems that combine the goals of stability and popularity. There was a group led by Bettina Klaus on lexicographic preferences in matching and market design. Finally, Florian Brandl led a group on the intersection between matching and fair division.

On the last day, there was a panel discussion that was moderated by Haris Aziz and Bettina Klaus. The discussants in the panel were Péter Biró, Ágnes Cseh, Lars Ehlers, Alex Teytelboym, and Utku Ünver. The main topics discussed included ways to build synergies between research communities and having an impact on the practice of matching markets.

The organisers thank all the Dagstuhl staff members for their professional support, the participants for enriching the seminar, Somouaoga Bonkoungou and Alex Lam for providing video conferencing support, and Seçkin Özbilen for his support in putting together the abstracts that compose the report.

6.20 Approximate Systems

Organizers: Eva Darulova, Babak Falsafi, Andreas Gerstlauer, and Phillip Stanley-Marbell

Seminar No. 21302

Date: July 25–30, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.6.147

Participants: Hussam Amrouch, David Atienza Alonso, Eric Atkinson, Andreas Burg, Eva Darulova, Lara Dolecek, Babak Falsafi, Djordje Jevdjic, Debasmita Lohar, Jürgen Teich, Damien Zufferey

Remote Participants: Sara Achour, R.Iris Bahar, Swarnendu Biswas, Peter Dueben, Andreas Gerstlauer, Ghayoor Gillani, Jie Han, Anastasiia Izycheva, Vijay Janapa Reddi, Gauri Joshi, Ulrich Kremer, Sasa Misailovic, Laura Monroe, Sri Parameswaran, Adrian Sampson, Olivier Sentieys, Phillip Stanley-Marbell, Radha Venkatagiri, Norbert Wehn, Georgios Zervakis

Resource efficiency is becoming an increasingly important challenge, especially due to the pervasiveness of computing systems and the diminishing returns from performance improvements of process technology scaling. At the same time, many important applications have nondeterministic specifications or are robust to noise in their execution. They thus do not necessarily require fully reliable computing systems and their resource consumption can be reduced by introducing or exposing approximations.

While trading correctness for efficiency has been part of computing systems since the early days, it has seen renewed interest in the past decade. Different techniques have been since developed for applying and controlling approximations and the errors they introduce at different levels of the compute stack. Unfortunately, most of these techniques have been applied in isolation, making simplified assumptions about the other levels. It is thus unclear how all the different techniques interact, combine and complement or negate each other to provide end-to-end application benefits.

The aim of this seminar was to bring together researchers from different domains working on approximate computing, algorithms, programming languages, compilers, architecture and circuits, in order to explore open challenges and opportunities and to define cross-area research directions and collaborations relating to an end-to-end application of approximate computing principles across the compute stack.

The seminar consisted of brief presentations by a subset of the participants that covered the entire computing stack from hardware to applications, and that focused on the current challenges. The talks were followed by discussions in breakout groups that first focused on the different application areas of high-performance computing, embedded systems and deep learning, followed by group discussions on particular fundamental and cross-cutting challenges that were identified during the first breakout session. The full report includes the abstracts of the participant's presentations as well as summaries of the breakout group discussions.

6.21 Coalition Formation Games

Organizers: Edith Elkind, Judy Goldsmith, Anja Rey, and Jörg Rothe

Seminar No. 21331

Date: August 15–20, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.7.1

Participants: Niclas Boehmer, Grégory Bonnet, Florian Brandl, Robert Bredereck, Martin Bullinger, Edith Elkind, Piotr Faliszewski, Shiri Heffetz, Martin Hoefer, Anna Maria Kerkmann, Bettina Klaus, Jérôme Lang, Christian Laußmann, Seckin Özbilen, Jörg Rothe, Sanjukta Roy

Remote Participants: Chris Addington, Nathan Arnold, Haris Aziz, Vittorio Bilo, Andreas Darmann, Gabrielle Demange, Hendrik Fichtenberger, Abhek Ghosh, Judy Goldsmith, Sushmita Gupta, Paul Harrenstein, Ayumi Igarashi, Joanna Kaczmarek, Panagiotis Kanellopoulos, Michael McKay, Tomasz P. Michalak, Anja Rey, Jacob Schlueter, Matthew Spradling, Taiki Todo, Anaëlle Wilczynski, Gerhard J. Woeginger, Makoto Yokoo, Yair Zick

As mentioned, coalition formation games occur in many real-world settings. We are particularly interested in a subclass of coalition formation games, hedonic games, which were first proposed by Drèze and Greenberg [1] and later formalized by Banerjee et al. [2] and Bogomolnaia and Jackson [3]. Hedonic games are distinguished from general coalition formation games by the requirement that each agent's utility is wholly derived from the members of their own coalition.

This Dagstuhl Seminar brought multiple approaches and viewpoints to the study of coalition formation games, and in particular hedonic games, mainly from the perspective of computer science and economics. Particular topics that were discussed in talks and working groups include:

- succinctly representable preferences over coalitions;
- evolving preferences;
- the existence and verification of stable coalition structures (for various stability concepts);
- the computational complexity of finding or verifying stable or optimal partitions, or even determining whether such partitions exist;
- designing (if possible, efficient) algorithms for finding stable or optimal (or nearly so) coalition structures, or for verifying that a coalition structure is (nearly) stable or optimal;
- stability notions restricted to social networks or other networks;
- matching markets and matching under preferences, and their relation to hedonic games;
- dynamics of coalition formation;
- and group activity selection.

The overarching theme of this Dagstuhl Seminar was to bring together different communities working in coalition formation and hedonic games from various perspectives in computer science and economics and to bridge and bundle their research activities.

Much of the great atmosphere of the seminars at Schloss Dagstuhl comes from informal meetings besides the official schedule, with participants doing leisure activities together and enjoying other joint undertakings – this is, by the way, coalition formation in practice. Owing to the hybrid mode and pandemic-related restrictions, it was unfortunately not possible for us to organize group activities with all participants. However, due to the great technical support at Schloss Dagstuhl, the participants – online and on site – were able to take part in talks, discussions and working groups interactively to explore some of the challenging open questions of the field.

The organizers thank all participants for interesting talks and discussions. We also thank Schloss Dagstuhl for the technical preparation and support that made this hybrid seminar possible.

References

1 Dreze, Jacques H and Greenberg, Joseph. *Hedonic Coalitions: Optimality and Stability*. Econometrica: Journal of the Econometric Society (1980): 987-1003.

2 Banerjee, Suryapratim, Hideo Konishi, and Tayfun Sönmez. *Core in a Simple Coalition Formation Game*. Social Choice and Welfare 18.1 (2001): 135-153.

3 Bogomolnaia, Anna, and Matthew O. Jackson. *The Stability of Hedonic Coalition Structures*. Games and Economic Behavior 38.2 (2002): 201-230.

6.22 Understanding I/O Behavior in Scientific and Data-Intensive Computing

Organizers: Philip Carns, Julian Kunkel, Kathryn Mohror, and Martin Schulz

Seminar No. 21332

Date: August 15–20, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.7.16

Participants: Wolfgang Frings, Yi Ju, Andreas Knüpfer, Julian Kunkel, Erwin Laure, Radita Liem, Frank Mueller, Sarah Neuwirth, Sebastian Oeste, Martin Schulz

Remote Participants: Marcus Vincent Boden, Jim Brandt, André Brinkmann, Suren Byna, Philip Carns, Fahim Tahmid Chowdhury, Hariharan Devarajan, Ann Gentile, Sivalingam Karthee, Roland Laifer, Jay Lofstead, Johann Lombardi, Stefano Markidis, Sandra Adriana Mendez, Kathryn Mohror, Michael Ott, Marc Snir, Shane Snyder, Mehmet Soysal, Osamu Tatebe, Devesh Tiwari, Chen Wang, Michèle Weiland, Weikuan Yu

Dagstuhl Seminar 21332, "Understanding I/O behavior in scientific and data-intensive computing," brought together computer scientists from around the world to survey how I/O workloads are measured and analyzed on high-performance computing (HPC) systems, identify gaps in methodologies, and debate how to best apply this technology to advance HPC productivity. The hybrid, week-long event attracted 10 physical and 25 virtual attendees. They included representatives from seven countries spanning a variety of career levels in academia, industry, and government. The diversity of perspectives, combined with an intense week-long seminar format, offered an unprecedented opportunity for researchers to share ideas and spark new collaborative opportunities.

The seminar agenda was structured as a combination of full-group plenary sessions and subgroup breakout sessions. The plenary sessions were used to discuss high-level issues, vote on subtopics to investigate, relay results from breakout sessions, and present "lightning" talks that highlighted key issues in the community. The breakout sessions employed small groups (roughly five people each) to follow up in "deep dive" discussions on specific subtopics. This format enabled attendees from numerous time zones to remain productively engaged throughout the week. We also found it to be successful in facilitating discussion despite the COVID-19 safety considerations that prevented us from assembling at a single venue. The final day of the seminar was devoted to recording seminar findings in a timely manner while subject matter experts were still available for consultation.

Over the course of the seminar, the attendees converged on six high-level topics for deep dive discussions that are covered in the full report.

- **Tools: Cross-Cutting Issues** (Section 4.1 of the full report) explored common challenges in development of tools for understanding HPC I/O.
- **Data Sources and Acquisition** (Section 4.2 of the full report) addressed how to acquire various forms of raw I/O instrumentation from production systems.
- **Analysis** (Section 4.3 of the full report) focused on how to interpret I/O instrumentation once acquired.
- **Enacting Actionable Responses** (Section 4.4 of the full report) investigated how to best utilize the outcomes from I/O analysis.
- **Data Center Support** (Section 4.5 of the full report) focused on strategies for facility operators to facilitate better understanding of I/O behavior.
- **Community Support** (Section 4.6 of the full report) explored the unique characteristics of the I/O analysis community and how to foster its growth.

The full report presents a separate summary for each deep dive topic, including a survey of the state of the art, gaps, challenges, and recommendations. The report concludes in Section 5 of the full report with a summary of cross-cutting themes and recommendations produced by the seminar as a whole. We found that understanding I/O behavior in scientific and data-intensive computing is increasingly important in an era of evolving workloads and increasingly complex HPC systems and that several cross-cutting challenges must be addressed in order to maximize its potential.

6.23 Identifying Key Enablers in Edge Intelligence

Organizers: Aaron Ding, Ella Peltonen, Sasu Tarkoma, and Lars Wolf

Seminar No. 21342

Date: August 22–25, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.7.76

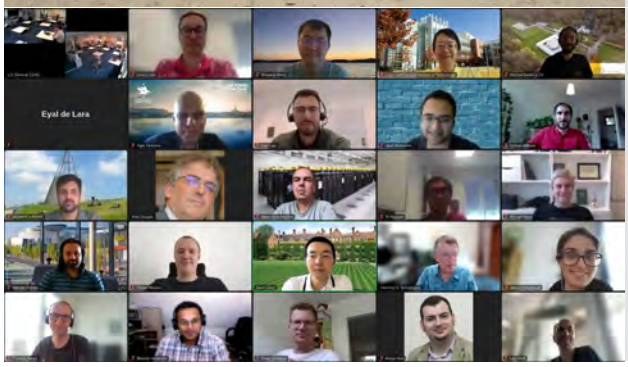

Participants: Christian Becker, Schahram Dustdar, Janick Edinger, Lauri Lovén, Tri Nguyen, Ella Peltonen, Jan Rellermeyer, Martijn Warnier, Lars Wolf

Remote Participants: Atakan Aral, Jari Arkko, Yiran Chen, Eyal De Lara, Aaron Ding, Fred Douglis, Diego Ferran, Thomas Hiessl, Dewant Katare, Dieter Kranzlmüller, Ling Liu, Madhusanka Liyanage, Ivan Lujic, Setareh Maghsudi, Nitinder Mohan, Iqbal Mohomed, Roberto Morabito, Petteri Nurmi, Jörg Ott, Francesco Regazzoni, Olga Saukh, Stefan Schulte, Henning Schulzrinne, Maarten Sierhuis, Stephan Sigg, Pieter Simoens, Gürkan Solmaz, Sasu Tarkoma, Wiebke Toussaint, Antero Vainio, Marten Van Dijk, Maarten van Steen, Blesson Varghese, Shiqiang Wang, Klaus Wehrle, Michael Welzl, Chenren Xu

Research Area

Edge computing, a key part of the upcoming 5G mobile networks and future 6G technologies, promises to decentralize cloud applications while providing more bandwidth and reducing latencies. The promises are delivered by moving application-specific computations between the cloud, the data-producing devices, and the network infrastructure components at the edges of wireless and fixed networks. The previous works have shown that edge computing devices are capable of executing computing tasks with high energy efficiency, and when combined with comparable computing power to server computers.

In stark contrast to the current edge-computing development, current artificial intelligence (AI) and in particular machine-learning (ML) methods assume computations are conducted in a powerful computational infrastructure, such as a homogeneous cloud with ample computational and data storage resources available. This model requires transmitting data from end-user devices to the cloud, requiring significant bandwidth and suffering from latency. Bringing computation close to the end-user devices would be essential for reducing latency and ensuring real-time response for applications and services. Currently, however, these benefits cannot be achieved as the perspective of "edge for AI", or even "communication for AI", has been understudied. Indeed, previous studies address AI only limitedly in different perspectives of the Internet of Things, edge computing, and networks.

Clear benefits can be identified from the interplay of ML/AI and edge computing. We divide this interplay into edge computing for AI and AI for edge computing. Distributed AI functionality can further be divided into edge computing for communication, platform control, security, privacy, and application or service-specific aspects. Edge computing for AI centres on the challenge of adapting the current centralized ML and autonomous decision-making algorithms to the intermittent connectivity and the distributed nature of edge computing. AI for edge computing, on the other hand, concentrates on using AI methods to improve the edge applications or the functionalities provided by the edge computing platform by enhancing connectivity, network orchestration, edge platform management, privacy or security, or providing autonomy and personalized intelligence on application level.

Previous studies address accommodating AI methods for different perspectives of IoT, edge computing and networks. However, there is still a need to understand the holistic view of AI methods and capabilities in the context of edge computing, com-

prising for example predictive data analysis, machine learning, reasoning, and autonomous agents with learning and cognitive capabilities. Further, the edge environment with its opportunistic nature, intermittent connectivity, and interplay of numerous stakeholders present a unique environment for deploying such applications based on computations units with different degrees of intelligence capabilities.

The AI methods used in edge computing can be further divided into learning and decision making. Learning refers to building, maintaining and making predictions with ML models, especially neural networks. Decision making is the business logic, that is, the process of acting upon the predictions. This is the domain of decision theory, control theory and game theory, whose solutions and equilibrium are now often estimated with data by reinforcement learning methods.

Currently, AI's cloud-centric architecture requires transmitting raw data from the end-user devices to the cloud, introducing latencies, endangering privacy and consuming significant data transmission resources. The next step, currently under active research, is distributed or federated AI, which builds and maintains a central model in the cloud or on the edge but allows user devices to update the model and use it locally for predictions. We envision a fully decentralized AI which flattens the distributed hierarchy, with the joint model built and maintained by devices, edge nodes and cloud nodes with equal responsibility. The present challenges for AI in edge computing converge on 1) finding novel neural network architectures and their topological splits, with the associated training and inference algorithms with fast and reliable accuracy and 2) distributing and decentralized model building and sharing into the edge, by allowing local, fast-to-build personalized models and global, collaborative models, and information sharing. Finally, the novel methods need to be 3) integrated with key algorithmic solutions to be utilised in edge-native AI applications. The ground-breaking objectives and novel concepts edge-native artificial intelligence brings are:

- Edge-native AI can be used for obtaining higher quality data from massive Internet of Things, Web of Things, and other edge networks by filtering out large volumes of noise, context labelling, dynamic sampling, data cleaning, etc. High-quality data can thus be used to feed both edge inferencing and cloud-based data analysis systems, for example, training large-scale machine learning models.
- Edge computing provides low latency that is crucial especially for real-time applications, such as anything related to driving and smart mobility. AI applications on the edge and thus closer to the end-user will not only fasten existing applications but also provide opportunities for novel and completely new solutions.
- Edge-based computing provides data privacy when users are involved, and no need to share data to the cloud services but only the locally learned model.
- With edge-computing implemented for AI/ML model building, personalisation of such models can be done in local environments without unnecessary transmission overhead (when only local data is anyway considered for model building). Global models built in the cloud environment can be used to support these local models whenever a collaborative, large or more general model is requested.
- Edge-native AI/ML tasks provide mobility of the computation and cloudlet-like processing in the edge. In comparison to cloudlets, edge computing provides more flexibility and dynamic operations for load balancing, task management, distribution of the models, etc.
- Light-weight computation on the edge devices and local environments can enable energy savings.
- Ethical data management: edge-native AI can be used to keep data ownership control closer to the user, e.g., when computation is managed and task distribution controlled from the user's own devices, and suitable security and privacy protection methods are in use.

6.24 Universals of Linguistic Idiosyncrasy in Multilingual Computational Linguistics

Organizers: Timothy Baldwin, William Croft, Joakim Nivre, and Agata Savary

Seminar No. 21351

Date: August 30–31, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.7.89

Remote Participants: Timothy Baldwin, Verginica Barbu Mititelu, Emily M. Bender, Archna Bhatia, Bernd Bohnet, Francis Bond, Cem Bozsahin, Ryan Cotterell, William Croft, Miryam de Lhoneux, Marie-Catherine de Marneffe, Jamie Findlay, Daniel Flickinger, Kim Gerdes, Voula Giouli, Tunga Gungor, Jan Hajic, Dag Haug, Uxoa Iñurrieta, Laura Kallmeyer, Christo Kirov, Maria Koptjevskaja Tamm, Artur Kulmizev, Lori Levin, Natalia Levshina, Teresa Lynn, Stella Markantonatou, Nurit Melnik, Paola Merlo, Yusuke Miyao, Kadri Muischnek, Joakim Nivre, Petya Osenova, Stephen Pepper, James Pustejovsky, Alexandre Rademaker, Carlos Ramisch, Manfred Sailer, Agata Savary, Emmanuel Schang, Nathan Schneider, Ivelina Stoyanova, Sara Stymne, Reut Tsarfaty, Francis M. Tyers, Meagan Vigus, Aline Villavicencio, Veronika Vincze, Ekaterina Vylomova, Nianwen Xue, David Yarowsky, Amir Zeldes, Daniel Zeman, Tim Zingler

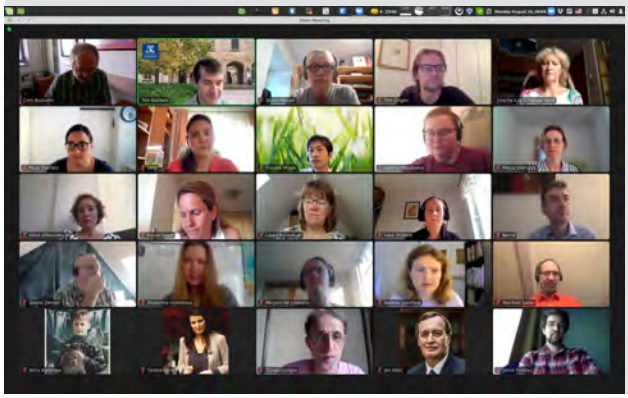

This Dagstuhl Seminar was initially planned as a 1-week event in June 2020 (with number 20261) with the following objectives:

- **Theoretical**: To deepen the understanding of language universals, and of how they apply to linguistic idiosyncrasy, so as to further promote unified modelling while preserving diversity.
- **Practical**: To improve the treatment of idiosyncrasy in treebanking frameworks, in computationally tractable ways and, thus, to foster high quality NLP tools for more languages with greater typological diversity.
- **Networking**: To promote a higher degree of convergence across typology-driven initiatives, while focusing on three main aspects of language modelling: morphology, syntax, and semantics.

Due to the COVID-19 pandemic, the event was first rescheduled and finally reduced to a 2-day online event on 30-31 August 2021, with two 3-hour sessions, repeated for better inclusiveness of various time zones (which corresponds to about 20% of the initially planned duration).

Prior to the event, participants submitted discussion issues, based on which working groups and the program were formed, as described in our Wiki space[42].

More precisely, the program of the event followed the Dagstuhl model:

- A list of recommended **readings** was published prior to the event
- **Introductory talks**, given by the 4 organizers, ensured common understanding of the scope and challenges to address.
- **Personal introductions** of all participants helped achieve a community building effect, despite the online setting.
- **Working groups** (WGs) were built on the basis of the discussion issues submitted by the participants. Each WG

[42] https://gitlab.com/unlid/dagstuhl-seminar/-/wikis/home

had 4 co-leaders, at least one of which could attend repeated sessions, so as to ensure consistency between the 2 time-zone sub-groups. The following WGs were created:
  - WG1: What counts as a word?
  - WG2: What counts as a MWE and as a construction?
  - WG3: Syntax vs. semantics
- **Discussion issues** were addressed in WGs by the proposers' short introductions followed by brainstorming.
- Plenary **reporting** sessions from WGs took place twice for every time zone.

The event attracted 51 participants, who judged it successful and expressed the need for a full-size on-site follow-up event. All the organizational details and outcomes of the seminar are gathered in our Wiki space[43].

Despite its very reduced and fully online format, the seminar achieved part of its objectives, stressed the importance of some initially-defined research questions, gave rise to new questions, and showed the efficiency of some instruments.

- On the **networking** side, the intended convergence effect was clearly apparent. While the initial proposal and invitee list was dominated by NLP-oriented members of the UD and PARSEME communities, strong contributions came notably from the less numerous typology and UniMorph experts. The four communities interacted actively, and reinforcing these interactions is intended for the near future. Notably, steps were taken towards:
  - integrating typology experts in the PARSEME core group
  - accompanying a seminal work in typology (Croft, to appear) with a "companion volume" about practical implementation of morphosyntactic concepts in UD.
- On the **theoretical** side, the event showed:
  - The importance of the research question *How to identify words across languages?* (item I.A in the seminar proposal), to which the whole of Working Group 1 was dedicated. In particular, new insights from lesser-studied languages, brought by typology experts, allowed us to broaden the perspective on this issue.
  - The need for capturing the relationship between the two fundamental notions in this proposal: a multiword expression and a construction, studied by Working Group 2. From the linguistic and typology perspective, a MWE is a special case of a construction, which is rarely made explicit in current NLP models. But the notion of a construction needs a more formal definition to be implementable in NLP, notably as far as the type-token opposition is concerned (question II.B in the seminar proposal). Thus, the typology-NLP interactions are essential in the quest for an optimal model.
  - The scope of the syntax-semantics interface issues (question II in the proposal) addressed by Working Group 3. On the one hand, the interests of the community in this respect exceeded the scope intended by the event organizers. Namely corpus-lexicon interlinking for all language units, not only for MWEs, was targeted. On the other hand, MWEs are exemplars of condensed syntax-semantic interface issues, and as such provide good case studies in this domain.
- On the **practical** side, some initial proposals emerged as to harmonizing UD treebank annotation guidelines with: (i) modelling morphological properties at the subword level (heavily studied by UniMorph), (ii) labelling MWEs (core activity of PARSEME).

Each multidisciplinary approach like ours bears heavy risks of intractability. This is because different communities often have different objectives and points of view on the same phenomena, and they may fail to agree on a unified approach, or even on the usefulness of working towards such a unification. In our case, there is a tension between:
- diversity and descriptive detail required in linguistics,
- necessary simplifications for the sake of robustness in NLP.

In other words, it is legitimate to question the usefulness of universality-driven initiatives (in NLP) if idiosyncrasy and diversity are basic properties of language data. Yet even typologists seek language universals which abstract away from the idiosyncrasy.

We feel that the event allowed us to mitigate this tension. Namely, even if a universality-based treebank fails to render the diversity of possible analyses of a language phenomenon, it is still useful not only for NLP applications but also for linguistic and typological analyses. This is because relevant examples are easy to extract (and to further re-interpret), as long as the annotation is consistent and well-documented.

Another barrier-lifting effect of the event concerned the relation between UD and PARSEME. It seems that the MWE categories defined by UD and PARSEME are less incompatible than initially expected, simply because the definition of an MWE in itself is different in UD and PARSEME. This could have been a source of major incompatibility but since a MWE does not really have a status in the UD annotation process, the discrepancies could (at least in some cases) be overcome relatively easily.

In conclusion, the event provided, in our opinion, a proof of concept for the framing objectives set up in the original Dagstuhl Seminar proposal. However, since the effective framework and duration was severely reduced as compared to the initially intended setting, only part of these objectives could be achieved. Thus, we are currently putting efforts to ensure follow-up events. In particular, a new Dagstuhl Seminar with roughly the same objectives has been submitted.

[43] https://gitlab.com/unlid/dagstuhl-seminar/-/wikis

6.25 Higher-Order Graph Models: From Theoretical Foundations to Machine Learning

Organizers: Tina Eliassi-Rad, Vito Latora, Martin Rosvall, and Ingo Scholtes

Seminar No. 21352

Date: August 29 to September 1, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.7.139

Participants: Unai Alvarez-Rodriguez, Luca Gallo, David F. Gleich, Christoph Gote, Jürgen Hackl, Andreas Hotho, Kang-Ju Lee, Leonie Neuhäuser, Vincenzo Perri, Giovanni Petri, Luka Petrovic, Martin Rosvall, Michael Schaub, Maximilian Schich, Ingo Scholtes, Frank Schweitzer, Markus Strohmaier

Remote Participants: Federico Battiston, Ginestra Bianconi, Rebekka Burkholz, Giulia Cencetti, Gabriele DiBona, Daniel Edler, Tina Eliassi-Rad, Anton Eriksson, Stephan Günnemann, Heather Harrington, Desmond J. Higham, Fariba Karimi, Danai Koutra, Renaud Lambiotte, Timothy LaRock, Vito Latora, Yamir Moreno, Natasa Przulj, Lisi Qarkaxhija, Alice Schwarze, Jelena Smiljanic, Michele Starnini, Chester Tan, Leo Torres, Anatol Wegner

The network science and graph mining community has created a rich portfolio of data analysis and visualisation techniques that have become a cornerstone for knowledge extraction from relational data on complex systems. Most of those techniques build on simple graph abstractions, where nodes represent a system's elements, and links represent *dyadic interactions, relations, or dependencies* between those elements. This mathematical formalism has proven useful for reasoning, e.g., about the centrality of nodes, the evolution and control of dynamical processes, and the community or cluster structure in complex systems, given that we have access to *relational data* [17]. However, the graph abstractions used in those methods typically do not account for **higher-order relations** between nodes that are present in many real complex systems. Important examples for such data include:

- relational data that is inherently non-dyadic, such as (unordered) sets of authors co-authoring scientific articles, protein triplets in a cell that simultaneously interact with each other, or actors in social systems engaging in group collaborations,
- time-stamped data on social networks with chronologically ordered sequences of (dyadic) interactions, where specific sequences of nodes interact via *causal paths*
- sequential data on networked systems, such as user click streams, mobility trajectories, financial transaction sequences, citation paths, or directed acyclic graphs that give rise to a chronologically or topologically ordered sequences of nodes traversed by processes
- data on networked systems with multiple types or layers of links that cannot be reduced to a simple graph model

Over the past years, researchers have shown that the presence of such higher-order interactions can fundamentally alter our understanding of complex systems. They can change our notion of the importance of nodes captured by centrality measures, affect the detection of cluster and community structures in graphs, and influence dynamical processes like diffusion or epidemic spreading, as well as associated control strategies in non-trivial ways [5, 10, 24, 27, 28, 33, 35]. To further develop graph-based representations of data and broaden their potential application in pattern recognition, data analysis, and machine learning, over the past few years researchers have developed a rich portfolio of **higher-order network models and representations** that capture more than just dyadic dependencies in complex systems. The organisers of this seminar have recently summarised current research and open challenges in this area in three independent overview and perspective articles [1, 15, 30]. An incomplete list of approaches explored over the past few years include:

- hypergraphs, where each *hyperedge* can connect an arbitrary number of nodes [11]
- simplicial network models, where *simplices* represent d-dimensional group interactions [10, 12]
- d-dimensional De Bruijn graphs, where edges capture *ordered* sets of d dyadic interactions [16, 28]

- memory networks, where memory nodes capture Non-Markovian properties in time series data [24]
- higher-, variable-, and multi-order Markov models for temporal networks [20, 27, 33]
- multi-layer and multiplex networks with multiple types of links between nodes [13]
- applications of categorical sequence mining techniques to model patterns in sequences of node sets [7]

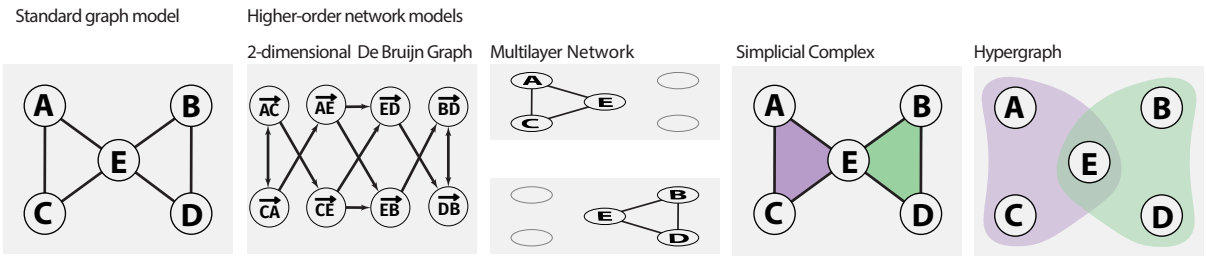

Fig. 6.12
Illustration of standard graph model (left) and four modelling approaches capturing different types of higher-order interactions proposed in topological data analysis, network science, and computer science. Figure adapted from [15].

In Figure 6.12, we illustrate some of the higher-order graph models listed above. All these modelling approaches address the same fundamental limitation of graph models when studying complex systems: **we cannot understand a system's structure and dynamics by decomposing direct and indirect interactions between elements into a set of dyadic relations with a single type**. However, the similarities and differences between these different approaches are still not fully understood.

At a critical time for the community, this Dagstuhl Seminar intended to improve our understanding of the strengths, weaknesses, commonalities, and differences of these different approaches along with their resulting computational and epistemological challenges. The seminar aimed to create a common foundation for developing graph mining and machine learning techniques that use recent advances in the study of higher-order graph models by gathering key researchers from different communities, including machine learning, information retrieval and data mining, complex systems theory, theoretical physics, network science, computational social science, and mathematics. The participants included senior and junior researchers focusing on four related and intersecting topics: (i) Topological and Graph-Theoretic Foundations, (ii) Higher-Order Models for Dynamical Processes, (iii) Higher-Order Pattern Recognition and Machine Learning, (iv) Computational Aspects in Higher-Order Graph Analysis and Graph Mining.

The organisers used the four topics to structure the seminar program and derive the participants' initial assignment to possible working groups. After an initial round of brief opening statements, participants introduced themselves and stated their specific interests for the seminar during five-minute lightning talks. During a match-making session taking place in the afternoon of day one, all interests expressed by the participants were consolidated into a set of working groups, addressing the following six areas: (i) Visualisation and Interpretability of Higher-Order Graph Models, (ii) Learning and Model Selection, (iii) Unification of Different Higher-Order Modelling Frameworks, (iv) Benchmark Data and Evaluation Practices, (v) Applications of Higher-Order Graph Models, and (vi) Societal Impact, Robustness, and Fairness. In the remaining time of the seminar, participants worked on those issues in the groups. The full report includes summaries of the opening statements, the results of the working groups, and a summary of a panel discussion taking place on the evening of day two.

References

1 Federico Battiston, Giulia Cencetti, Iacopo Iacopini, Vito Latora, Maxime Lucas, Alice Patania, Jean-Gabriel Young, and Giovanni Petri. Networks beyond pairwise interactions: structure and dynamics. *Physics Reports*, 2020.

2 Austin R. Benson. Tools for higher-order network analysis. *CoRR*, abs/1802.06820, 2018.

3 Austin R. Benson, Rediet Abebe, Michael T. Schaub, Ali Jadbabaie, and Jon Kleinberg. Simplicial closure and higher-order link prediction. *Proceedings of the National Academy of Sciences*, 115(48):E11221–E11230, 2018.

4 Austin R. Benson, David F. Gleich, and Jure Leskovec. Tensor spectral clustering for partitioning higher-order network structures. In *Proceedings of the 2015 SIAM International Conference on Data Mining*, pages 118–126, 2015.

5 Austin R Benson, David F Gleich, and Jure Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016.

6 Austin R. Benson, David F. Gleich, and Lek-Heng Lim. The spacey random walk: A stochastic process for higher-order data. *SIAM Review*, 59(2):321–345, 2017.

7 Austin R Benson, Ravi Kumar, and Andrew Tomkins. Sequences of sets. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1148–1157, 2018.

8 Matthias Bolten, Stephanie Friedhoff, Andreas Frommer, Matthias Heming, and Karsten Kahl. Algebraic multigrid methods for laplacians of graphs. *Linear Algebra and its Applications*, 434(11):2225 – 2243, 2011. Special Issue: Devoted to the 2nd NASC 08 Conference in Nanjing (NSC).

9 Daniel Edler, Ludvig Bohlin, et al. Mapping higher-order network flows in memory and multilayer networks with infomap. *Algorithms*, 10(4):112, 2017.

10 Ernesto Estrada and Grant J. Ross. Centralities in simplicial complexes. applications to protein interaction networks. *Journal of Theoretical Biology*, 438:46 – 60, 2018.

11 Gourab Ghoshal, Vinko Zlatić, Guido Caldarelli, and Mark EJ Newman. Random hypergraphs and their applications. *Physical Review E*, 79(6):066118, 2009.

12 Iacopo Iacopini, Giovanni Petri, Alain Barrat, and Vito Latora. Simplicial models of social contagion. *Nature communications*, 10(1):1–9, 2019.

13 Mikko Kivelä, Alex Arenas, Marc Barthelemy, James P. Gleeson, Yamir Moreno, and Mason A. Porter. Multilayer networks. *Journal of Complex Networks*, 2(3):203–271, 07 2014.

14 Renaud Lambiotte, Martin Rosvall, Michael Schaub, Ingo Scholtes, and Jian Xu. Beyond graph mining: Higher-order data analytics for temporal network data. In *Hands-On Tutorial at the 24th ACM SIGKDD International Conference on Knowledge Discovery*, volume 38, 2018.

15 Renaud Lambiotte, Martin Rosvall, and Ingo Scholtes. From networks to optimal higher-order models of complex systems. In *Nature physics*, Vol. 15, p. 313-320, March 25 2019

16 Timothy LaRock, Vahan Nanumyan, Ingo Scholtes, Giona Casiraghi, Tina Eliassi-Rad, and Frank Schweitzer. HYPA: Efficient detection of path anomalies in time series data on networks. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 460–468. SIAM, 2020.

17 Vito Latora, Vincenzo Nicosia, and Giovanni Russo. *Complex Networks: Principles, Methods and Applications*. Cambridge University Press, 2017.

18 Naoki Masuda, Mason A. Porter, and Renaud Lambiotte. Random walks and diffusion on networks. *Physics Reports*, 716-717:1 – 58, 2017. Random walks and diffusion on networks.

19 Huda Nassar, Caitlin Kennedy, Shweta Jain, Austin R. Benson, and David F. Gleich. Using cliques with higher-order spectral embeddings improves graph visualizations. In *Proceedings of the 2020 Web Conference (WWW)*, pages 2927–2933, 2020.

20 Tiago P Peixoto and Martin Rosvall. Modelling sequences and temporal networks with dynamic community structures. *Nature communications*, 8(1):1–12, 2017.

21 Vincenzo Perri and Ingo Scholtes. Hotvis: Higher-order time-aware visualisation of dynamic graphs. In *Proceedings of the 28th International Symposium on Graph Drawing and Network Visualization (Graph Drawing 2020), Vancouver, BC, Canada (to appear)*, 2020.

22 Luka V. Petrovic and Ingo Scholtes. PaCo: Fast Counting of Causal Paths in Temporal Network Data. In *Proceedings of the 11th Temporal Web Analytics Workshop (TempWeb 2021)* held in conjunction with The Web Conference 2021, Ljubljana, Slovenia, April 2021

23 Luka V. Petrovic and Ingo Scholtes. Learning the markov order of paths in a network, 2020. arXiv 2007.02861.

24 Martin Rosvall, Alcides V Esquivel, Andrea Lancichinetti, Jevin D West, and Renaud Lambiotte. Memory in network flows and its effects on spreading dynamics and community detection. *Nature communications*, 5:4630, 2014.

25 Mandana Saebi, Giovanni Luca Ciampaglia, Lance M Kaplan, and Nitesh V Chawla. Honem: Network embedding using higher-order patterns in sequential data. *arXiv preprint arXiv:1908.05387*, 2019.

26 Mandana Saebi, Jian Xu, Lance M. Kaplan, Bruno Ribeiro, and Nitesh V. Chawla. Efficient modeling of higher-order dependencies in networks: from algorithm to application for anomaly detection. *EPJ Data Sci.*, 9(1):15, 2020.

27 Ingo Scholtes. When is a network a network? multi-order graphical model selection in pathways and temporal networks. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '17, page 1037–1046, New York, NY, USA, 2017. Association for Computing Machinery.

28 Ingo Scholtes, Nicolas Wider, René Pfitzner, Antonios Garas, Claudio Tessone, and Frank Schweitzer. Causality-driven slow-down and speed-up of diffusion in non-markovian temporal networks. *Nature communications*, 5:5024, 2014.

29 Yingxia Shao, Shiyue Huang, Xupeng Miao, Bin Cui, and Lei Chen. Memory-aware framework for efficient second-order random walk on large graphs. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, page 1797–1812, 2020.

30 Leo Torres, Ann S Blevins, Danielle S Bassett, and Tina Eliassi-Rad. The why, how, and when of representations for complex systems. *arXiv preprint arXiv:2006.02870*, 2020.

31 Francesco Tudisco, Austin R. Benson, and Konstantin Prokopchik. Nonlinear higher-order label spreading. *CoRR*, abs/2006.04762, 2020.

32 Tao Wu, Austin R. Benson, and David F. Gleich. General tensor spectral co-clustering for higher-order data. In *Advances in Neural Information Processing Systems 29*, pages 2559–2567, 2016.

33 Jian Xu, Thanuka L Wickramarathne, and Nitesh V Chawla. Representing higher-order dependencies in networks. *Science advances*, 2(5):e1600028, 2016.

34 Hao Yin, Austin R. Benson, Jure Leskovec, and David F. Gleich. Local higher-order graph clustering. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 555–564, 2017.

35 Yan Zhang, Antonios Garas, and Ingo Scholtes. Higher-order models capture changes in controllability of temporal networks. In Journal of Physics: Complexity, Vol. 2, No. 1, January 29, 2021

6.26 Extending the Synergies Between SAT and Description Logics

Organizers: Joao Marques-Silva, Rafael Peñaloza, and Uli Sattler

Seminar No. 21361

Date: September 5–10, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.8.1

Participants: Meghyn Bienvenu, Maria Paola Bonacina, Stefan Borgwardt, Camille Bourgaux, David Carral, Aaron Eberhart, Jean-Marie Lagniez, Joao Marques-Silva, Jordi Planes, Anni-Yasmin Turhan

Remote Participants: Carlos Ansotegui, Leopoldo Bertossi, Olaf Beyersdorff, Emmanuel Hebrard, Ian Horrocks, Alexey Ignatiev, Matti Järvisalo, Oliver Kullmann, Carlos Mencía, Thomas Meyer, Jakob Nordström, Magdalena Ortiz, Ana Ozaki, Rafael Peñaloza, Francesco Ricca, Uli Sattler, Renate Schmidt, Stefan Szeider, David Tena Cucala, Cesare Tinelli, Renata Wassermann

About the Seminar

Propositional satisfiability (SAT) and Description Logics (DL) are two successful areas of computational logic where automated reasoning plays a fundamental role. Seen from a very abstract level, they can be thought as being part of the same family of logical formalisms attempting to represent knowledge from an application domain, and differentiated only by their expressivity and correspondent trade-off in reasoning complexity. However, the evolution of the two areas has diverged, mainly due to differences in their underlying goals and methods. While the DL community focused on introducing and fully understanding new constructors capable of expressing different facets of knowledge, the SAT community built highly-optimised solvers targeted for industrial-size problems.

Some recent work has permeated the boundaries between the two communities. It has been shown that some DL reasoning problems could be reduced to known SAT-related tasks. In turn, these reductions motivated new optimisations targeted to the specific shape of the problems constructed by them. The goal of this seminar was to bring together researchers from both communities to foster a deeper collaboration and mutual development. The primary goals were (i) to understand the tasks and methods from one community which could benefit the other, and (ii) to discuss the policies used within the communities to encourage specific advancements.

A relevant issue considered is how to promote the development and testing of DL reasoners. To try to answer this, we discussed the status of benchmarks in both communities, and the success stories from SAT competitions. A salient point was the issue, from the DL point of view, of the many variants that should be evaluated – from the different languages, to the reasoning tasks considered. However, recent SAT competitions have also successfully handled many categories. One possible explanation for the wide availability of solvers capable of handling practical extensions of SAT (like MaxSAT and QBF) is the existence of solvers like MiniSAT, which allow for fast prototyping using SAT solvers as oracles. No analogous tool is available for DL reasoners.

The remaining of the seminar focused on novel and timely tasks which are currently under development in both communities, and where the best possibilities for collaborations are foreseen. Among them, we can mention methods for explaining the result from a solver, and proofs which can be used to automatically verify their correctness. We noted that the notion of an *explanation* is too wide, allowing for different interpretations which were presented as talks during the seminar. Each of these interpretations gives rise to distinct techniques. But interestingly, the core ideas are not necessarily specific to SAT or DLs. This last observation can lead to collaborations studying the problems from both points of view.

In addition to the longer talks whose abstracts accompany this document, other impromptu presentations were triggered by the previous discussions. One clear conclusion which can be taken from these engagements is that the potential for synergic growth between the areas is large and worth exploring.

Format

Due to the COVID-19 situation, the seminar had to be held in a hybrid format. While this had the obvious disadvantage of limiting the social interactions and offline scientific discussions that characterise Dagstuhl Seminars, it also allowed the participation of many who, by distance or travel limitations, would have not been able to attend.

Overall, the hybrid format meant having a more structured and linear program than originally planned for the seminar, but as mentioned already the results are promising.

6.27 Structure and Learning

Organizers: Tiansi Dong, Achim Rettinger, Jie Tang, Barbara Tversky, and Frank van Harmelen

Seminar No. 21362

Date: September 5–10, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.8.11

Participants: Tiansi Dong, Thomas Liebig, Alexander Mehler, Sven Naumann, Achim Rettinger, Steffen Staab, Volker Tresp

Remote Participants: Mehwish Alam, Christian Bauckhage, Sven Behnke, Jian-Jia Chen, Anthony Cohn, Armin B. Cremers, Bernado Cuenca Grau, Asja Fischer, Xiaoming Fu, Jürgen Gall, Christian Hempelmann, Mateja Jamnik, Markus Knauff, Michael Kramer, Jens Lehmann, Juanzi Li, Pietro Lio, Yong Liu, Zhiyuan Liu, Xiaoxing Ma, Tristan Miller, Siba Mohsen, Adithya Murali, Roberto Navigli, Max Petrenko, Julia Rayz, Ron Sun, Jian Tang, Jie Tang, Barbara Tversky, Frank van Harmelen, Ran Yu

Deep Learning systems are the hope of the fifth industrial revolution. However, recent studies have found that Deep Learning systems can be easily manipulated, i.e. in Natural Language Understanding, Object Recognition. How to introduce structures into Deep Learning systems to improve reliability and performance has become a hot topic in Natural Language Processing (NLP), Machine Learning (ML), Semantic Web (SW) communities around the world. The aim of the seminar is to bring together interdisciplinary researchers around the world for constructive discussions on this theme, in particular, it intends to establish international collaborations to promote computational Humor, with the hope to let AI bring more joy, more laugh into the world, and do more good for the society. The hybrid seminar is structured in the form of Talks, Working Groups, and Open questions. The seminar started with the talk "Hybrid AI for Humor". The dynamic semantics of humor is beyond the reach of the classic symbolic AI, the deep learning paradigm, and current neural-symbolic integration methods, but can be captured by the neural geometric embedding, in terms of rotating sphere embedding. This novel embedding is rooted in Qualitative Spatial Representation (QSR) in symbolic AI and Learning Representation (LR) in neural ML. The former tries to symbolically delineate the basic spatial knowledge that humans have and possible ways that this knowledge can be used as a reference for abstract knowledge in other domains. LR aims at learning latent feature knowledge from data. The motivation and a geometric approach to realizing the unification were introduced in the talk "Rotating Spheres – A New Wheel for Neuro-Symbolic Unification". The motion of rotating spheres in high-dimensional space is served as a computational model to simulate (1) the motion of the physical world, (2) the circular interaction among the mind, the body, and the world (called *spraction* – a contraction of space, action, and abstraction, in which actions in space create abstractions).

The motion of the physical world is vividly explained in the talk "Rotating Spheres in the Milky Way". This *spraction* process is explained in the talk "Thinking with the Body and the World", which can guide the design of novel cognitive robots, and promote novel cognitive architectures. Two topics were covered by the talk "Learning about Language and Action for Robots", and the talk "Neural-Symbolic Models, Dual-Process Theories, and Cognitive Architectures".

In primates, the same brain structures that support spatial thinking also support conceptual thinking, Single cells in hippocampus gather multi-media information from different memories in the brain to represent places in space, events in time, ideas in conceptual spaces. Update-to-date research of neural simulation is introduced by Volker Tresp with the talk "Knowledge Graph and Cognitive Learning: from Perception to Memory Embedding", which maps embedding models to various cognitive memory functions, in particular to semantic and concept memory, episodic memory, sensory memory, short-term memory, and working memory.

Spatial thinking is multi-modal and established and distorted by our actions and perceptions of the spaces we interact in. This raises two questions: What are good representations for video understanding? and how to compute symbolic rules that the models have learned from the training data? Juergen Gall introduced holistic video understanding and argued the potential of hybrid approaches that combine neural networks with symbolic AI for video understanding and reasoning. Cuenca Grau, Bernado gave the talk "Characterizing Graph Neural Networks Using Logical Rules". He formally defines what it means for a set of logical rules to characterize the behavior of a model and proposes a GNN-based architecture that admits a characterization in terms of Datalog rules.

Spatial thinking is evident in the ways we think and the ways we externalize thought, for example, through words. Our words act on thought the way we act on objects. The philosophy of spatial thinking challenges the computational approach to natural language processing and understanding. Roberto Navigli argued that Natural Language Understanding (NLU) is particularly challenging, as this requires the machine to go beyond processing strings to reach a semantic level. Recent developments and challenges were discussed through three key tasks in NLU, namely Word Sense Disambiguation, Semantic Role Labeling, and Semantic Parsing. Zhiyuan Liu argued that knowledge (including symbols, embeddings, or models) is the key to a deeper understanding of human languages and that big pretrained language models can be regarded as the most advanced approach to model knowledge and to capture knowledge (including commonsense) from plain text and that the key challenge is how to incorporate both open data and structural knowledge. Alexander Mehler reviewed problems of neural network-based language learning, suggested to introduce the concept of cognitive maps and spatial information processing, and sketched a synergistic model that relates the dynamics of distributed information processing to bias interaction. Jie Tang introduced Wu-Dao, China's first homegrown super-scale intelligent model system, with the goal of building an ultra-large-scale cognitive-oriented pretraining model to focus on essential problems in general artificial intelligence from a cognitive perspective. Wu-Dao substantially outperforms BERT on the SuperGLUE natural language understanding benchmark with the same amount of pre-training data. Alam Mehwish discusses the characteristics of the existing benchmark datasets for the task of KG Completion, and limitations of the existing benchmark datasets and targets those issues in the generation of LiterallyWikidata.

Another externalization of spatial thinking is through graphics. In the talk "Semi-Riemannian Graph Convolutional Networks", Steffen Staab introduced their new geodesic tools that allow for extending neural network operations into geodesically disconnected semi-Riemannian manifolds. Thomas Liebig introduced using p-adic coding and computation for structured domains or domains with inherent granularity.

The ultimate form of spatial thinking is comics (a form of humor, the most creative form of storytelling), which typically show bodies acting in space. Humor is used as a testbed and lighthouse for the development of AI and machine learning. In the talk "Ethics of AI Humor" Kiki explained how humor has frustrated symbolic and statistic AI approaches; in the talk "Knowledge and Inferences Needed for Humor" Julia Rayz introduced recent advances in transformer-based approaches, and raised open questions.

Working groups are the main components of the seminar. The hybrid seminar provides an excellent chance to practice the situation that participants can continue to work together after this seminar, which is the main outcome of this seminar.

The seminar ended with the discussion "Boxology for Hybrid Learning and Reasoning Systems" chaired by Frank van Harmelen.

6.28 Integrated Deduction

Organizers: Maria Paola Bonacina, Philipp Rümmer, and Renate A. Schmidt

Seminar No. 21371

Date: September 12–17, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.8.35

Participants: Erika Abraham, Alexander Bentkamp, Jasmin Christian Blanchette, Maria Paola Bonacina, Pascal Fontaine, Carsten Fuhs, Stéphane Graham-Lengrand, Ullrich Hustadt, Moa Johansson, Patrick Koopmann, Magnus Myreen, Ruzica Piskac, Philipp Rümmer, Renate Schmidt, Claudia Schon, Stephan Schulz, Alexander Steen, Martin Suda, Sophie Tourret, Sarah Winkler

The full report contains the program and outcomes of the Dagstuhl Seminar 21371 on *Integrated Deduction* that was held at Schloss Dagstuhl, Leibniz Center for Informatics, during September 12–17, 2017. It was the fourteenth in a series of Dagstuhl Deduction seminars held biennially since 1993.

The motivation for this seminar was the following. Automated deduction has developed a wide and diverse range of methods and tools for logico-deductive reasoning. They include SAT solvers,[44] SMT solvers,[45] automated theorem provers, aka ATP systems, proof assistants, aka interactive theorem provers (ITP), as well as libraries of formalized mathematics and formalized knowledge. These methods and tools have found successful application in computing fields as diverse as the *analysis, verification, and synthesis of systems*, *programming language design*, *knowledge engineering*, and *computer mathematics*. However, no method, tool, paradigm, or even reasoning style can solve all problems, or respond to all demands coming from even a single field of application. Therefore, the next grand challenge for automated deduction is *integration*.

Integration occurs and is needed at different abstraction levels. Within deduction itself, integration of deductive engines allows us to build more powerful, more flexible, more expressive reasoners, that can solve more problems with fewer resources, meaning not only memory and computing time, but also human time and human expertise, the latter two often being the most precious of resources. Next, deductive reasoners get integrated into other tools, such as *automated test generators*, *verifying compilers*, or *program synthesizers*, just to name a few. Yet another level of integration occurs when *logico-deductive reasoning* is integrated with other forms of automated reasoning, such as *probabilistic reasoning* and *statistical inference*. This leads to the integration of deduction within *intelligent systems*, such as *decision support systems*, *agent programming environments*, and *data processing systems*. Here deduction may provide explanation, course of action, and the capability of learning from missing information; it may also aid modelling and facilitate agent communication.

The seminar on *Integrated Deduction* successfully covered as many as possible of these integration issues, including:

- Integration of deductive engines into more general automated deductive systems;
- Integration of automated deductive systems into interactive proof assistants;
- Integration of deduction into formal methods tools;
- Integration of deduction for knowledge processing; and
- Integration of deduction into intelligent systems such as agent-based systems.

Furthemore, the seminar investigated a number of key technological and human-related issues, that are largely orthogonal to most integration contexts, affecting both feasibility and deployment of integrated deduction. Examples of such issues are:

- The development of interfaces for integration;
- The generation of continuous feedback during the run of deductive tools, including also information from intermediate or unsuccessful states;
- The reproducibility of results in the presence of tool updates or imposed resource limits (e.g., available computation time or memory) that may introduce non-determinism; and
- Advanced tradeoff's between performance and expressivity as well as between specialization and genericity.

Practical challenges around integrated deductive systems, includ-

[44] SAT solvers are solvers for satisfiability queries in propositional logic, known as the SAT problem.
[45] SMT stands for satisfiability modulo theories.

ing collaboration with non-expert users or access to data sets, were also discussed.

The seminar brought together a diverse audience, including both researchers working in deduction and researchers working in neighbouring areas that make use of deduction. Many participants have experience in building, using, and applying systems with integrated deduction.

Following the tradition of the Dagstuhl *Seminars on Deduction*, most of the program consisted of contributed talks by participants on their research. In this manner, the bottom-up style of the Dagstuhl experience was preserved, allowing for spontaneous contributions as they emerged during the seminar.

However, this seminar was also innovative with respect to tradition, in that it featured *five invited tutorials* on key topics in integrated deduction. These tutorials were valuable in highlighting the state-of-the-art in the integration of deduction systems and in fomenting discussions on challenges and open problems.

The program also featured a hike in the woods and a social dinner in a nearby village, that helped establishing or strengthtening collaborations.

The following section contains the abstracts for most of the talks and tutorials listed in alphabetical order.

6.29 Behavioural Types: Bridging Theory and Practice

Organizers: Mariangiola Dezani, Roland Kuhn, Sam Lindley, and Alceste Scalas

Seminar No. 21372

Date: September 12–17, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.8.52

Participants: Marco Carbone, Simon Fowler, Philipp Haller, Mathias Jakobsen, Eduard Kamburjan, Roland Kuhn, Sam Lindley, Fabrizio Montesi, Philip Munksgaard, Alceste Scalas, Peter Thiemann, Emilio Tuosto

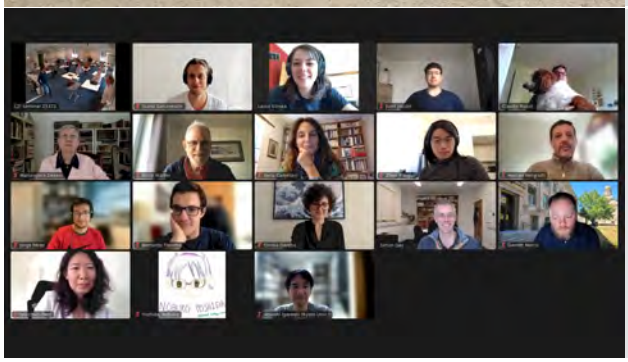

Remote Participants: Gul Agha, Stephanie Balzer, Christian Bartolo Burló, Laura Bocchi, Ilaria Castellani, Tzu-Chun Chen, Ornela Dardha, Mariangiola Dezani, Simon Gay, Raymond Hu, Atsushi Igarashi, Jules Jacobs, Wen Kokke, Dimitrios Kouzapas, Hernán Melgratti, J. Garrett Morris, Rumyana Neykova, Marco Peressotti, Jorge A. Pérez, Claudio Russo, Guido Salvaneschi, Bernardo Toninho, Vasco T. Vasconcelos, Laura Voinea, Philip Wadler, Nobuko Yoshida, Fangyi Zhou

This seminar followed the earlier Dagstuhl Seminar 17051 "Theory and Applications of Behavioural Types". Whereas Seminar 17051 was quite broad, encompassing both theory and practice across a wide range of areas relating to behavioural types, this seminar was much more focused, concentrating on how best to enable the use of behavioural types for practical programming.

Initial preparations

We gathered initial lists of proposed talks and breakout topics prior to the start of the seminar via an online form. We added to these throughout the week. We scheduled talks and breakout groups daily depending on audience interest and participant availability. The first part of the week was primarily talks, with ample time for stimulating discussions; the second part included more time for breakout sessions.

Hybrid seminar logistics

Due to the SARS-CoV-2 pandemic, the seminar was organised in hybrid format, with both in-person and remote participants. As the virtual participants came from a wide range of time zones (from central US to Japan) we gave special consideration to the time slot 2pm–4pm CEST during which everyone could attend. Those in Europe and Japan were able to attend morning sessions and those in Europe and America to attend further afternoon sessions (and a special evening session on Monday).

In order to run a successful hybrid Dagstuhl Seminar, we made essential use of the dedicated equipment available at Dagstuhl: a Zoom-based streaming setup, with multiple cameras and ceiling microphones in the seminar room. All talks were live-streamed to both virtual and in-person participants. Talks were recorded so that virtual participants from incompatible time zones could catch up, then deleted at the end of the week. Larger hybrid breakout sessions were held in the main seminar room, and smaller ones elsewhere using a more ad hoc setup.

Moreover, all participants (local and virtual) were invited to use Zulip (a chat application) to exchange messages and files, pose questions during presentations, and remain informed on the upcoming events, group activities, and schedule updates.

Activities and outcomes

Throughout the seminar, the participants gathered in focused breakout groups: the findings of the breakout groups are described in more detail elsewhere in the report. Here is a brief summary:
Typing non-channel-based models allowed researchers with a

wide range of perspectives and backgrounds to exchange their views. A key observation was that modern concurrent systems that coordinate via streams of events are difficult to analyse and verify with using existing approaches, and new formalisms are needed.

Logic-based approaches reviewed the state of the art, and discussed new directions. One of the conclusions is that more research is needed to relate concurrent and distributed systems to a broader range of logics beyond classical and intuitionistic linear logic (which are the focus of most current publications).

Type-informed recovery strategies explored failure handling at different levels (from network to application), and summarised several open questions not addressed in existing work.

Session types with untrusted counter-parties focused on how to ensure that different processes interact under compatible protocols, establishing the beginning of new work on monitoring and adaptation.

Join patterns / synchronisation – the next generation collected a survey of various attempts to integrate join patterns in programming languages, and discussed why they have not yet become mainstream. The discussion highlighted the need for exploring the connections between join patterns and linear logic, and the use of the join calculus as a reference for new implementation attempts.

The participants of several breakout groups have agreed to continue their work and collaboration after the seminar.

In addition to these more structured breakout sessions there were further lively improvised meetings and discussions (especially after dinner) which are not summarised in the report.

Overall, we believe that the seminar activities were a success. Unfortunately the hybrid format did pose a barrier for remote participants, especially those in different time zones. But on the positive side, for many participants this was their first Dagstuhl Seminar, and for the in-person participants it was their first in-person scientific gathering after many months of virtual events due to the SARS-CoV-2 pandemic: their feedback has been enthusiastic.

At the end of the seminar the participants agreed to remain in contact to continue the discussions, and foster new collaborations. There was strong enthusiasm for organising a follow-up Dagstuhl Seminar in the future, perhaps taking place in about two years time. To enable future collaborations the participants:

1. created a GitHub organisation where all seminar participants (and other researchers invited later) can exchange references and materials;
2. agreed to use the seminar's Zulip chat (mentioned above) as a starting point to set up a more permanent solution for continuing the interactions and exchanges (e.g., a mailing list);
3. nominated four people who will propose a new seminar, building upon the results of this one.

6.30 Conversational Agent as Trustworthy Autonomous System (Trust-CA)

Organizers: Effie Lai-Chong Law, Asbjørn Følstad, Jonathan Grudin, and Björn Schuller

Seminar No. 21381

Date: September 19–24, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.8.76

Participants: Elisabeth André, Oliver Bendel, Leigh Clark, Asbjørn Følstad, Frode Guribye, Sebastian Hobert, Andreas Kilian, Dimosthenis Kontogiorgos, Matthias Kraus, Guy Laban, Effie Lai-Chong Law, Minha Lee, Clayton Lewis, Birthe Nesset, Catherine Pelachaud, Martin Porcheron, Stefan Schaffer, Ryan Schuetzler, Björn Schuller, Eren Yildiz

Remote Participants: Theo Araujo, Susan Brennan, Heloisa Candello, Ana Paula Chaves, Cristina Conati, Benjamin Cowan, Laurence Devillers, Jasper Feine, Jonathan Grudin, Evelien Heyselaar, Soomin Kim, Stefan Kopp, Yi-Chieh Lee, Oliver Lemon, Q. Vera Liao, Christine Liebrecht, Roger K. Moore, Stefan Morana, Cosmin Munteanu, Ana Paiva, Symeon Papadopoulos, Caroline Peters, Rolf Pfister, Olivier Pietquin, Aleksandra Przegalinska, Elayne Ruane, Marita Skjuve, Cameron Taylor, Ricardo Usbeck, Margot van der Goot, Dakuo Wang, Saskia Wita, Levi Witbaard, Zhou Yu, Michelle X. Zhou

The overall goal of the Dagstuhl Seminar 21381 "Conversational Agent as Trustworthy Autonomous System" (Trust-CA) was to bring together researchers and practitioners, who are currently engaged in diverse communities related to Conversational Agents (CA), to explore challenges in maximising the trustworthiness of and trust in conversational agents as AI-driven autonomous systems – an issue deemed increasingly significant given their widespread uses in every sector of life – and to chart a roadmap for the future conversational agent research. The three main challenges we identified were:

- How do we develop trustworthy conversational agents?
- How do we build people's trust in them?
- How do we optimise human and conversational agent collaboration?

The Seminar Trust-CA took place on 19-24 September 2021 in a hybrid mode. Out of 50 invitees, 19 attended in person and the rest joined online from all over the world, including Brazil, Canada, France, Germany, Greece, Ireland, Netherlands, Norway, Poland, South Korea, Sweden, Switzerland, UK and USA.

The four-day scientific programme started by unpacking the notion of "trust in conversational agent" with a panel discussion. Each of the four seminar organisers expressed their views on the notion. Jonathan Grudin presented a list of ten species of trust that can be applied to conversational agents, for instance, "Trust that a CA will correctly interpret my question or request; will deliver relevant, reliable, useful information." Asbjørn Følstad first presented an overview of the six themes derived from a pre-Seminar survey (details are in Overview of Working Groups) and then described his recent work on the effect of human likeness of a conversational agent on trust. Björn Schuller presented factors influencing trust in humans, such as being reliable, ethical, moral and charismatic, and in conversational agents, such as being explainable, interpretable and transparent. He also discussed how to measure trust reliably and the danger of overtrust. Effie Law discussed the notion of trust with reference to multidisciplinary theory of trust (e.g., psychological, social, historical), beyond the use of questionnaires to evaluate trust, and identifying applications where agents are of high practical value. Some attendees commented on the ideas shared, e.g., the elusiveness of trust.

The scientific programme comprised two major parts – Talks and Breakout Groups. There were altogether 20 talks, covering a range of topics (see Abstracts). Nine of the talks were delivered in person and the rest online. There were six Breakout Groups with each discussing one of the six themes: Group 1 – Scope of Trust in CA; Group 2 – Impact of CA; Group 3 – Ethics

of CA; Group 4 – AI and Technical Development; Group 5 – Definition, Conceptualisation and Measurement of Trust; Group 6 – Interaction Design of CA. Group 1, 3 and 4 had one team each whereas Group 2, 5 and 6 had two teams each. To ease collaboration, individual teams were either in-person or online (except for Group 4 which was in hybrid mode). Each group had three two-hour working sessions. In the evening, each group reported progress and invited feedback for shaping subsequent sessions.

The group discussions led to intriguing insights that contributed to addressing the main challenges listed above and stimulated future collaborations (see the Workgroup Reports). Here we highlight one key insight of each group. Group 1 developed a dynamic model of trust with three stages, Build-Maintain-Repair, which evolve over time. Group 2 drafted a code of ethics for trustworthy conversational agents with eight provisions. Group 3 explored the ethics challenge of transparency from the perspective of conversational disclosure. Group 4 called for increased collaboration across research communities and industries to strengthen the technological basis for trust in conversational agents. Group 5 proposed a framework for integrating measurement of trusting beliefs and trusting behaviour. Group 6 analysed several aspects of multimodality to understand their possible effects on trust in conversational agents. Apart from the scientific programme, the Seminar organised several social events, including after-dinner wine and cheese gatherings, hiking in a nearby historic site, and a music event.

Overall, our Dagstuhl Seminar Trust-CA was considered a success. The major outputs were derived from the pre-Seminar survey (six research themes and a recommended reading list), twenty talks, and six multi-session breakout groups. Thanks must go to the enthusiastic involvement of all attendees in analysing various aspects of the burgeoning topic of conversational agents. Of course, the Seminar could only take place with the generosity of *Schloss Dagstuhl – Leibniz Center for Informatics*. The efficiency and friendliness of the scientific and administrative staff of Schloss Dagstuhl was much appreciated by the organisers and all attendees.

6.31 Sparsity in Algorithms, Combinatorics and Logic

Organizers: Daniel Král', Michał Pilipczuk, Sebastian Siebertz, and Blair D. Sullivan

Seminar No. 21391

Date: September 26 to October 1, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.8.115

Participants: Marthe Bonamy, Édouard Bonnet, Marcin Brianski, Jan Dreier, Zdenek Dvorák, Kord Eickmeyer, Fedor V. Fomin, Jakub Gajarský, Martin Grohe, Meike Hatzel, Petr Hlinený, Gwenaël Joret, Noleen Koehler, Martin Koutecký, Daniel Král', Stephan Kreutzer, Nikolas Mählmann, Piotr Micek, Wojciech Nadara, Jaroslav Nešetril, Jan Obdržálek, Patrice Ossona de Mendez, Kristýna Pekárková, Marcin Pilipczuk, Michal Pilipczuk, Felix Schröder, Sebastian Siebertz, Dimitrios M. Thilikos, Szymon Torunczyk, Torsten Ueckerdt, Alexandre Vigny, Bartosz Walczak

Remote Participants: Saeed Amiri, Daniel Cranston, Anuj Dawar, Peter Gartland, Archontia C. Giannopoulou, Mario Grobler, Shweta Jain, Ken-ichi Kawarabayashi, O-joung Kwon, Aurelie Lagoutte, Brian Lavallee, Rose McCarty, Yosuke Mizutani, Sang-il Oum, Filip Pokrývka, Evangelos Protopapas, Wojciech Przybyszewski, Daniel Quiroz, Felix Reidl, Benjamin Rossman, Saket Saurabh, Nicole Schirrmacher, Nicole Schweikardt, Luc Segoufin, Marek Sokolowski, Giannos Stamoulis, Blair D. Sullivan, Nate Veldt, David Wood

Motivation

It was realized already in the early days of computer science that structures (networks, databases, etc.) that are sparse appear ubiquitously in applications. The sparsity of input can be used in a variety of ways, e.g. to design efficient algorithms. This motivates a theoretical study of the abilities and limitations of sparsity-based methods. However, a priori it is not clear how to even define sparsity formally. Multiple sparsity-oriented paradigms have been studied in the literature, e.g. bounded maximum or average degree models, topologically constrained classes of graphs or graphs with bounded width parameters. However, many of those paradigms suffer from being either too restrictive to model real-life applications, or too general to yield strong tractability results.

In the late 2000's, Nešetřil and Ossona de Mendez proposed a new framework of uniform structural sparsity for classes of graphs that generalized existing definitions and initiated the development of a toolbox of sparsity-based methods for analyzing graphs. The central notions of their framework are bounded expansion and nowhere dense classes. It quickly turned out that the proposed notions can be used to build a mathematical theory of sparse graphs that offers a wealth of tools, leading to new techniques and powerful results. This theory has been extensively developed in the recent years.

It is particularly remarkable that the concepts of classes of bounded expansion and nowhere dense classes can be connected to fundamental ideas from multiple other fields of computer science, often in a surprising way, providing several complementary viewpoints on the subject. On one hand, foundations of the area are grounded in structural graph theory, which aims at describing structure in graphs through various decompositions and auxiliary parameters. On the other hand, nowhere denseness seems to delimit the border of algorithmic tractability of first-order logic, providing a link to finite model theory and its computational aspects. Finally, there is a fruitful transfer of ideas to and from the field of algorithm design: sparsity-based methods can be used to design new, efficient algorithms, especially in the paradigms of parameterized complexity, approximation algorithms, and distributed computing. Further, classic techniques for designing algorithms on sparse inputs inspire new combinatorial results on sparse graphs.

The aim of this Dagstuhl Seminar is to bring together researchers working on various aspects of sparsity in their own fields, in order to facilitate the exchange of ideas, methods and questions between different communities. So far, the synergy effect between graph theory, logic, and algorithm design has led to fundamental developments in the theory of sparse graph classes. Our goal is to inspire a new wave of developments by "stirring in the pot" of researchers working on different facets of sparsity.

An important part of the seminar will be the discussion of the (still fledgling) area of real-life applications of sparsity-based methods, where theory and practice could meet.

Seminar organization

Due to the on-going COVID pandemic, the seminar was held in hybrid format. In total the seminar was attended by 61 participants around the world (from North America to Europe to Asia). 32 of the participants were on-site and 29 remote. To make the hybrid format a success and in particular allow all members to participate in talks and working groups, the following measures were taken.

1. To accommodate both on-site and remote participants, a mix of on-site and remote talks were scheduled. The talks were scheduled in the early to late afternoon (MEZ local time), allowing remote audience members from all parts of the world to attend.
2. Both on-site and remote talks were streamed via zoom. The zoom session was projected onto a whiteboard in the seminar room. The remote participants could see and hear the on-site whiteboard and slide presentations. They could interrupt and ask questions or ask questions in the chat, which were then read by the organizers. This turned out to be a quite successful setup in which all participants could discuss in real-time.
3. On the first day of the seminar we had a short introduction of all participants, one invited tutorial lecture, one contributed talk and the open problem session. In total, we had 5 tutorial lectures and 12 contributed talks spread over the week. The topics and speakers were chosen to create a joint understanding of the state of the art in the fields that were brought together in the seminar.
4. The remaining program put a strong emphasis on open time for ad-hoc discussions and working in groups. After the open problem session on Monday, several groups of on-site and remote participants were formed, who approached the posed problems.
5. A discord server was set up to coordinate further communication and to keep track of the progress of the working groups.
6. A social event was organized online on Tuesday evening.

Work on open problems

Following the open problem session on the first day of the seminar, spontaneous groups working on selected open problems emerged. These typically included a mix of on-site and online participants, working in either synchronous or asynchronous manner using the Discord platform as a mean of communication. Below we list a selection of directions that were pursued during the seminar.

Model-checking on interpretations of locally well-behaved structures. It is known that model-checking First Order logic (FO) can be done in fixed-parameter time on classes of graphs that are locally well-behaved, for instance have locally bounded treewidth. However, the question is whether this is still true if the input graph is "logically disguised", or more precisely, has been additionally mapped through some FO transduction. The aim of this research group was to provide an affirmative answer by proving the following theorem: For every class of graphs $\mathcal{C}$ that is stable and can be transduced from a class of locally bounded cliquewidth, the model-checking problems for FO is fixed-parameter tractable on $\mathcal{C}$. This would generalize several known results on efficient FO model-checking on classes of dense graphs, e.g. map graphs or interpretations of classes of bounded degree, as well as provide multiple new results.

Transducing paths from classes of unbounded shrubdepth. The emerging logically-motivated structure theory for graphs uses First-Order transductions as the main notion of embedding. It is important to understand possible duality theorems for this notion, of the following form: If a class of graphs $\mathcal{C}$ does not admit a decomposition of some form, then $\mathcal{C}$ transduces a class of specific obstacles witnessing this conclusion. The aim of this research group was to prove the most basic conjecture following this pattern: If a class of graphs $\mathcal{C}$ has unbounded shrubdepth, then $\mathcal{C}$ transduces the class of all paths.

Treedepth vs pathwidth. It is known that every graph of pathwidth $\Omega(ab)$ has treewidth at least a or contains a binary tree of depth b as a minor. It is also known that every graph of treedepth $\Omega(abc)$ has treewidth at least a, or contains a binary tree of depth b as a minor, or contains a simple path of length 2^c. This suggests the following conjecture: every graph of treedepth $\Omega(bc)$ has either pathwidth at least b or contains a simple path of length 2^c. The aim of this research group was to resolve this conjecture.

Treewidth-twin-width. The definition of the recently introduced graph parameter twin-width revolves around the mechanism of contraction sequences: simplification operations using which one can "fold" the whole graph into a single vertex. The main complexity measure of a contraction sequence is the maximum number of error edges that are adjacent to any vertex at any time. The goal of this group was to investigate the combinatorial properties of a graph parameter dubbed *treewidth-twin-width* obtained by additionally requiring that at all times, the graph composed of the error edges has bounded treewidth. Of particular interest is whether various classes known to have bounded twin-width actually have bounded treewidth-twin-width.

Integer programs equivalent to ones with bounded primal treedepth. Integer programming is known to be efficiently solvable for instances with small primal or dual treedepth. While we have a relatively good understanding of conditions when the instance can be transformed to an instance with small dual treedepth, less is known in the case of primal treedepth. The aim of this research group was to relate, for a given instance of integer programming, the smallest possible primal treedepth of an equivalent instance to invariant properties of the instance itself, in particular, to the structure of the column matroid of the constraint matrix. The ultimate goal would be to design algorithms for constructing an equivalent instance with small primal treedepth while avoiding a blow up in the entry complexity (such a blow up would prevent using the existing IP algorithms to solve the constructed instance).

Acknowledgements

Schloss Dagstuhl provided an excellent environment for hosting the seminar. The seminar room was appropriate to host the on-site participants and we found plenty of room for continuing discussions and socializing outside of the official program. This is particularly remarkable in these difficult times with the ongoing COVID pandemic. All participants were eager to meet and research together. According to the conducted survey, as well as the informal feedback to the organizers, the seminar was highly appreciated and can be considered a full success. On behalf of all participants, the organizers want to express their gratitude to the entire Dagstuhl staff and their outstanding support and service throughout the seminar.

6.32 Visualization of Biological Data – From Analysis to Communication

Organizers: Karsten Klein, Georgeta Elisabeta Marai, Kay Katja Nieselt, and Blaz Zupan

Seminar No. 21401

Date: October 3–8, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.9.1

Participants: Jan Aerts, Emma Beauxis-Aussalet, Michael Behrisch, Katja Bühler, Nadezhda T. Doncheva, Mennatallah El-Assady, Carsten Görg, Theresa Anisja Harbig, Lynda Hardman, Lawrence Hunter, Helena Jambor, Andreas Kerren, Karsten Klein, Stephen G. Kobourov, Alexander Koch, Michael Krone, Martin Krzywinski, Alexander Lex, Lennart Martens, Torsten Möller, Kay Katja Nieselt, Bruno Pinaud, Timo Ropinski, Falk Schreiber, Cagatay Turkay, Blaz Zupan

Remote Participants: Jillian Aurisano, Tanya Berger-Wolf, Katy Börner, Andreas Bueckle, Guadalupe Canahuate, Jian Chen, Nils Gehlenborg, Barbora Kozlíková, Raghu Machiraju, Georgeta Elisabeta Marai, James Procter, William Ray, Jos B.T.M. Roerdink, Ryo Sakai, Hagit Shatkay, Charlotte Soneson, Miha Štajdohar, Marc Streit, Granger Sutton, Danielle Szafir, Mathias Witte Paz

Advances in technology have turned biology into data-driven research. High-throughput and high-resolution techniques help us generate and collect vast amounts of data to be explored, analyzed, and turned into knowledge or actionable models. This abundance of biological data creates a substantial challenge in processing, analysis, and modeling. A popular way to address this challenge is through visual representation and analysis of the data. Creating compelling visualizations of biological data requires combining data visualization, bioinformatics, statistics, and computational biology. The field of biological visualization is interdisciplinary and involves collaboration between researchers from different areas.

Our aim with this Dagstuhl Seminar was to bring together researchers from multiple disciplines to discuss how to continue the interdisciplinary dialogue and foster the development of an international community concerned with biological visualization. We aimed to examine the state of the art and find areas to advance the research that might benefit from the joint efforts of all groups involved.

Our initial aim, expressed in the seminar proposal, was to explore the following four topics:

- data abstraction to support building custom visual tools of biological data,
- interactive analysis for biological data exploration,
- collaboration and communication through new tools,
- curriculum for teaching visualization in bioinformatics.

We have discussed these topics in the first two days of our five-day seminar and gradually came out with the following six working groups:

- facilitating cross-expertise exploration in explainable AI for multi-omics via visualization,
- visions for the lab notebook of the future,
- visual analytics of multilayer networks representing knowledge graphs,
- recommendations for designing visual, interpretable, and deep learning-based analytics pipelines in medical imaging,
- semantically enabled biomedical cartooming,
- a curriculum for the future of biological visualization.

Notice that with this new set of topics for the working groups, our seminar still closely followed our initial aim to explore the space of data abstractions, interactive analysis, and design of tools to support collaborations.

We have developed the schedule for the seminar based on our experience and expertise in previous successful Dagstuhl Seminars. We aim to emphasize the balance between prepared talks and panels and breakout groups for less structured discussions focused on a selection of highly relevant topics. Three types of plenary presentations were available to participants who had indicated an interest in presenting during the seminar:

- Overview talks (20 minutes plus 10 minutes for questions)
- Regular talks (10 minutes plus 5 minutes for questions)
- Panel presentations (5 minutes per speaker followed by a 20–25 minute discussion)

The breakout groups met multiple times for several hours during the week and reported to the overall group on several occasions. This format successfully brought bioinformatics and visualization researchers onto the same platform and enabled researchers to reach a shared, deep understanding through their questions and answers. It also stimulated fruitful discussions that all participants deeply appreciated.

We have organized the seminar during the COVID-19 pandemic. Due to various regulations and quarantines, slightly above half of the participants attended in person, while the other participants attended online. The meeting took the hybrid form, and we thank Dagstuhl for equipping the seminar rooms with suitable hardware. Still, we found the organization of the hybrid meeting challenging, to say the least, as it imposed constraints on the discussion and engagement of everyone. An all-important part of Dagstuhl's experience is off-line meetings during meals, trips, or long evenings, which online participants miss. Thus, also following the responses in the participants' survey, we would endorse their recommendation that Dagstuhl should return to the previous, non-hybrid format of seminars once the pandemic stops.

The full report describes in detail the outcomes of our meeting. At the present stage, the outcome includes a set of white papers summarizing the breakout sessions, overviews of the talks, and an emerging detailed curriculum for future biological and medical visualization education.

Acknowledgements

We would like to thank all participants of the seminar for their contributions and lively discussions; we also would like to thank the scientific directorate of Dagstuhl for providing us with the opportunity to organize this seminar. Finally, the seminar would not have been possible without the untiring help of the (scientific) staff of Dagstuhl, including Ms. Jutka Gasiorowski and Ms. Susanne Bach-Bernhard.

6.33 Digital Disinformation: Taxonomy, Impact, Mitigation, and Regulation

Organizers: Claude Kirchner and Franziska Roesner

Seminar No. 21402

Date: October 3–6, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.9.28

Participants: Camille Darche, Lynda Hardman, Claude Kirchner, Jean-Yves Marion

Remote Participants: Esma Aimeur, Jos Baeten, Asia J. Biega, Sébastien Gambs, Krishna Gummadi, Vladimir Kropotov, Evangelos Markatos, Filippo Menczer, Trisha Meyer, Franziska Roesner, Kavé Salamatian, Juliette Sénéchal, Dimitrios Serpanos, Serena Villata

Dagstuhl Seminar 21402 on Digital Disinformation occurred on October 4–6, 2021. The seminar was initially planned by Claude Kirchner (CNPEN/CCNE & Inria), Ninja Marnau (CISPA), and Franziska Roesner (University of Washington), and it was then co-lead and the report was written by Kirchner and Roesner, with input from other seminar participants. The seminar had been originally planned for June of 2020 but was then postponed due to the COVID-19 pandemic. It was held in a hybrid format, with some participants on site in Dagstuhl and most others joining remotely via the video conferencing system Zoom.

In order to maximize discussion and allow the interests of the group to drive the direction of the seminar, we did not plan for formal talks. Participants were asked to prepare a single slide, few-minute introduction about their research interests and methodologies related to digital disinformation, and a "burning question" they have in the space.

Participants included the following individuals, spanning a range of expertise from computer science to law:

- Esma Aïmeur (University of Montréal, Canada)
- Jos Baeten (CWI Amsterdam, Netherlands)
- Asia Biega (Max Planck Institute for Security and Privacy, Germany)
- Camille Darche (CNPEN, Inria and Université Paris Nanterre, France)
- Sébastien Gambs (Université du Québec à Montréal, Canada)
- Krishna Gummadi (Max Planck Institute for Software Systems, Germany)
- Claude Kirchner (CNPEN/CCNE and Inria, France)
- Vladimir Kropotov (Trend Micro, Russia)
- Jean-Yves Marion (Lorraine University, France)
- Evangelos Markatos (University of Crete, Greece)
- Fil Menczer (Indiana University, USA)
- Trisha Meyer (Vrije Universiteit Brussel, Belgium)
- Franziska Roesner (University of Washington, USA)
- Kavé Salamatian (University of Savoie, France)
- Juliette Sénéchal (University of Lille, France)
- Dimitrios Serpanos (University of Patras, Greece)
- Serena Villata (CNRS, France)

Based on our preliminary discussions, we identified four topics of interest to many seminar participants: *trustworthiness algorithms* (i.e., how to build systems that assess trust automatically), *friction as a technique in platform design* (e.g., to allow for people to take time and a step back when consuming information on social media), *the ethics of interventions* (e.g., the ethics of blocking or content moderation), and *how to educate users* (e.g., without creating over-skepticism). We then structured the rest of the seminar around four deep-dive conversations on these topics, described in the subsequent sections of the full report. Due to the relatively small size of the gathering, and most participants' broad interest in all four topics, we did not break out into smaller discussion groups but rather continued to discuss as a full group.

6.34 Machine Learning in Sports

Organizers: Ulf Brefeld, Jesse Davis, Martin Lames, and Jim Little

Seminar No. 21411

Date: October 10–15, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.9.45

Participants: Gabriel Anzer, Arnold Baca, Pascal Bauer, Ulf Brefeld, Jesse Davis, Björn Eskofier, Dennis Faßmeyer, Eric Hayman, Arno J. Knobbe, Otto Kolbinger, Philipp Kornfeind, Martin Lames, Daniel Link, Björn Mäurer, Fabrizio Pece, Martin Rumo, Tiago Guedes Russomanno, Marc Schmid, Jan Van Haaren, Benedicte Vanwanseele, Albrecht Zimmermann

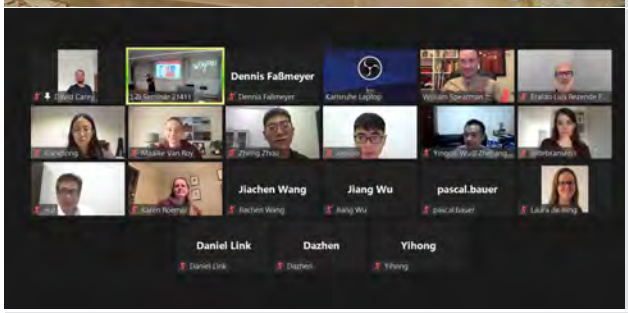

Remote Participants: Luke Bornn, Lotte Bransen, Mirjam Bruinsma, David Carey, Xiangtong Chu, Laura de Jong, Uwe Dick, James Elder, Irfan A. Essa, Mehrsan Javan, Jim Little, Patrick Lucey, Konstantinos Pelechrinis, Eraldo Luis Rezende Fernandes, Karen Roemer, Yannick Rudolph, Oliver Schulte, William Spearman, Karl Tuyls, Maaike Van Roy, Jiang Wu, Yingcai Wu, Hui Zhang, Zhou Zheng

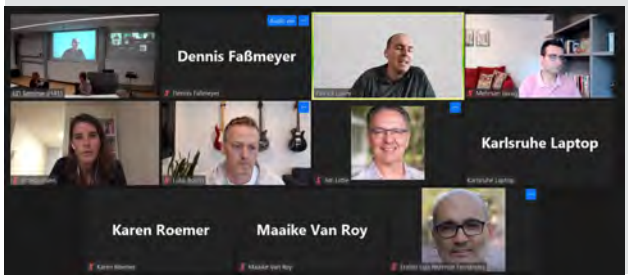

Sports has become an incredibly data rich field with the advent of data sources such as event data (e.g., time and locations of actions), tracking data (i.e., positional data), and athlete monitoring (e.g., bio-sensors, IMUs, GPS). These data are commonly and widely collected across multiple different sports, both on a professional and recreational level. The advent of such data raises the need to exploit the collected data both from the theoretical (e.g., sports modeling) as well as practical (e.g., training in top level sports) perspective. Problem-solving solutions can only be provided by an interaction between the sports science & informatics (S&I) and the machine learning (ML) communities. Machine learning is emerging as a powerful, new paradigm for sports analytics, as it provides novel approaches to making sense of the collected data. However, the S&I and ML communities are traditionally separate, each with its own agenda. The seminar aims to bring together top researchers and practitioners who are active in these two fields that can contribute to an assessment of their potential synergies.

We structured the seminar along five different themes, each of which was the focus of between half and a full day. Given the diversity of the participants' backgrounds in terms of discipline, each theme began with overview to get everyone on the same page. Then there were more detailed presentations. The five themes were:

Machine learning meets sports The goal of this session was to provide an overview of some of the machine learning techniques (predictive modeling, text mining) and how they can be applied in the sports context. The illustrative applica-

tions where ML can play a role included assessing the performances of teams and players, supporting sport broadcast, assessing fans reactions to rule changes and helping reduce the time burden on video analysts.

Sports science meets machine learning The goal was to provide an overview of basic concepts in sports science to inform researchers from machine learning. The basic concepts were the relation between competition, training and athlete's abilities, the structure of performance in different sports, and the demand for support in sports practice. In particular, the structure of team sports as dynamic interaction processes with emergent behavior was explained as this the most frequent application field for machine learning in sports.

Computer vision for sports The session aimed to expose the participants to the practice of gathering information about team sports through analysis of visual information. The session began with an overview of the general practice of computer vision for sports. Three of the presenters are from industry, representing companies with significant presence in the business of providing to analytics producers information on team sports such as basketball, football (soccer), and ice hockey. The fourth presenter, from academia, discussed material on camera planning and analytics and in addition has himself been involved in tech transfer of visual analytics methods from amateur sports. The overall goal, i.e., informing the participants regarding methods and applications of vision, was well met by the lectures of these experienced researchers.

Interdisciplinary view on tactics The session aimed to build a common understanding of tactics and their implementation in predictive/generative models. It is still an open question how to represent overarching long-term strategies in computer models and different ideas were discussed on the example of overview and contributed presentations.

Explaining, interpreting, and visualizing models and data for sports A key challenge is effectively conveying the results of machine learned models to domain experts, which is compounded by the black-box nature of many such models. This session highlighted a variety of techniques for meeting this objective, with illustrative examples arising from practice were shown for a variety of sports such as ice hockey, table tennis and football. This remains an active area of research and variety of lessons learned and ideas for improving the communication between domain-experts and technical-experts were discussed.

6.35 Quantum Cryptanalysis

Organizers: Stacey Jeffery, Michele Mosca, María Naya-Plasencia, and Rainer Steinwandt

Seminar No. 21421

Date: October 17–22, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.9.64

Participants: Marco Baldi, Jean-François Biasse, Xavier Bonnetain, André Chailloux, Thomas Debris-Alazard, Yfke Dulek, Martin Ekerå, Stacey Jeffery, Antoine Joux, Stavros Kousidis, Nils Gregor Leander, Frédéric Magniez, Maria Naya-Plasencia, Phong Q. Nguyen, Alexandru Paler, Galina Pass, Edoardo Persichetti, Stephanie Reinhardt, Paolo Santini, Claus Peter Schnorr, André Schrottenloher, Nicolas Sendrier, Yixin Shen, Jana Sotáková, Rainer Steinwandt, Jean-Pierre Tillich

Remote Participants: Andris Ambainis, Shi Bai, Aleksandrs Belovs, Daniel J. Bernstein, Jintai Ding, Philippe Gaborit, András Gilyén, Maria Isabel González Vasco, Akinori Hosoyamada, Tetsu Iwata, Samuel E. Jaques, Floyd Johnson, Elena Kirshanova, Péter Kutas, Tanja Lange, François Le Gall, Dustin Moody, Michele Mosca, Ludovic Perret, Rachel Player, Thomas Pöppelmann, Angela Robinson, Yu Sasaki, John M. Schanck, Daniel C. Smith-Tone, Fang Song, Adriana Suárez Corona, Dániel Szabó, Bo-Yin Yang

Motivation and scope

Owing to the ongoing pandemic, this (sixth) installment of the Dagstuhl Seminar series on *Quantum Cryptanalysis* was held in a hybrid format. The focus of this seminar was on deployed schemes and more mature post-quantum cryptographic schemes, such as Round 3 candidates in NIST's standardization effort. For the technical program of the seminar, we encouraged research on

Quantum algorithmic innovations to attack cryptographic building blocks, leveraging state-of-the-art quantum computing. How can we leverage quantum algorithms to improve cryptanalytic capabilities, and how can we optimize the best available cryptanalytic results in meaningful quantum attack models?

Techniques and software tools to optimize and quantify resources for such attacks. Can we establish reasonably precise quantum resource counts for cryptanalytic attacks, especially for problem instances and parameter choices that are actually deployed or considered for standardization for future deployment?

Quantum attacks against today's RSA or elliptic-curve based cryptography and against modern block ciphers, which help us understand the urgency for transitioning to post-quantum solutions, fall in the seminar scope. As in the past, the seminar brought together researchers who work in the field of quantum computing with experts in classical cryptography, taking into account the latest advances in both fields. With 26 participants on site and 29 remote participants, Schloss Dagstuhl hosted a broad group of leading experts from across the globe.

Organization

The ongoing pandemic impacted the organization of the seminar, which for the first time was offered in a hybrid format. Thanks to the available technology at Schloss Dagstuhl and the efficient support of two volunteers (Shaun Kepley and Galina Pass), integrating remote presentations into the schedule worked smoothly.

The scheduling accounted for time zone differences and, as in the past, we left ample time for discussions and collaboration – for a typical day, we scheduled no more than four presentations. Following the Dagstuhl tradition and in line with prior seminars in the Quantum Cryptanalysis series, there was no technical program during Wednesday afternoon, leaving participants time for exploring the surroundings, spending time on research, or taking care of testing requirements for upcoming international travel.

Results and next steps

The collaboration between cryptographers and experts in quantum computing has come a long way, and it seems fair to say that this Dagstuhl Seminar series has contributed to this positive development. The quantum cryptanalytic progress in symmetric cryptography is very noticeable. This was evidenced by the number and quality of presentations on this subject offered by seminar participants. On the asymmetric side, the presentations demonstrated fascinating research progress on understanding computational problems related to lattices and codes. At the same time, a need remains to better quantify the potential of quantum algorithms for tackling hardness assumptions as used in state-of-the-art post-quantum proposals.

6.36 Rigorous Methods for Smart Contracts

Organizers: Nikolaj S. Bjørner, Maria Christakis, Matteo Maffei, and Grigore Rosu

Seminar No. 21431

Date: October 24–29, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.9.80

Participants: Elvira Albert, Leonardo Alt, Nikolaj S. Bjørner, Maria Christakis, Marco Eilers, Grzegorz Fabianski, Josselin Feist, Bryan Ford, Diego Garbervetsky, Isabel Garcia-Contreras, Arthur Gervais, Neville Grech, Wolfgang Grieskamp, Fritz Henglein, Matteo Maffei, Alexander Nutz, Sophie Rain, Albert Rubio, Tanja Schindler, Yannis Smaragdakis, Valentin Wüstholz

Remote Participants: Massimo Bartoletti, Andreea Buterchi, Jing Chen, Shuo Chen, Ankush Das, Stefan Dziembowski, Ákos Hajdu, Aniket Kate, Markulf Kohlweiss, Igor Konnov, Yi Li, Victor Luchangco, Anastasia Mavridou, Noam Rinetzky, Grigore Rosu, Giulia Scaffino, Gerardo Schneider, Ilya Sergey, Zhong Shao, Philip Wadler, Yoni Zohar

The seminar attracted 22 on-site and approximately as many off-site participants. The hybrid mode presented an opportunity for collaborators, particularly students, of invitees to participate remotely and contribute to the discussions. Remote participation spanned all time zones which attested to their involvement. The on-site participants had the benefit of extended interactions and relation building so crucial for advancing scientific activities.

The technical program was organized around first day of tutorial presentations on the main topics covered by the seminar. These topics were *static analysis techniques*, *program verification methods*, *protocol design for decentralized ledgers*, and *semantic-based tools*.

The following days provided for in-depth sessions around these topics. Static analysis techniques spanned using Horn clause solvers, Datalog engines, and abstract interpretation frameworks in a mixture of academic and industrial settings. Program verification techniques, likewise, were pursued both by academic and industry participants. The seminar offered an excellent forum for the scientific and commercial community around smart contracts to exchange experiences and develop ideas.

For the social program, we hiked for two hours during a beautiful October afternoon to Landgasthof Paulus & Der Laden for a delightful dinner.

6.37 Probabilistic Numerical Methods – From Theory to Implementation

Organizers: Philipp Hennig, Ilse C.F. Ipsen, Maren Mahsereci, and Tim Sullivan

Seminar No. 21432

Date: October 24–29, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.9.102

Participants: Simon Bartels, Nathanael Bosch, François-Xavier Briol, Maurizio Filippone, Giacomo Garegnani, Roman Garnett, Alexandra Gessner, Philipp Hennig, Toni Karvonen, Peter Nicholas Krämer, Maren Mahsereci, Katharina Ott, Marvin Pförtner, Jonathan Schmidt, Tomas Teren, Filip Tronarp, Jonathan Wenger, Stephen Wright

Remote Participants: Oksana Chkrebtii, Jon Cockayne, Yuhan Ding, Matthew Fisher, Fred J. Hickernell, Nick Higham, Ilse C.F. Ipsen, Motonobu Kanagawa, Hans Kersting, Tadashi Matsumoto, Michael McKerns, Masha Naslidnyk, Chris Oates, Michael A. Osborne, Houman Owhadi, Andrei Paleyes, Kamran Pentland, Geoff Pleiss, Jagadeeswaran Rathinavel, Timothy Reid, Simo Särkkä, Florian Schäfer, Aleksei Sorokin, Tim Sullivan, Aretha Teckentrup, Onur Teymur, Zi Wang

Probabilistic Numerical algorithms frame a numerical task as a statistical inference problem, expressed in the language of probabilistic inference. The key advantage of this approach is that it allows quantification of uncertainty arising from finite computational resources, and to combine thus with other forms of uncertainty, in particular those arising from model misspecification, finite observational data, and measurement errors. In recent years, algorithms arising from this formalism have repeatedly shown that they can enrich and improve upon classic methods in tasks where

- hyperparameter adaptation is not straightforward;
- computational stochasticity and low precision play a prominent role;
- limited data make uncertainty quantification a key functionality;
- related problems have to be solved repeatedly;
- and where extreme scale or tight budgets call for rough approximations at low cost.

Probabilistic Numerics lies at the intersection of machine learning within computer science and numerical analysis within applied mathematics. This interdisciplinary nature raises an exciting and challenging set of viewpoints with regards to goals and challenges of the field. The first goal of this seminar was to rekindle our community following two years of pandemic lockdown, to provide an opportunity to update others on one's own research, and to discuss new directions and ideas together. We were lucky to assemble – both in-person and remote – a diverse group of people from computer science, machine learning, from statistics, optimization, and from numerical analysis.

The second key goal of this seminar was to take the next step in the development of probabilistic numerical methods by focusing on their implementation. From `Lapack` to `SciPy` to `PyTorch`, open-source software libraries have driven scientific advancement in their respective domains. Such libraries accelerate research, enable benchmarking and promote the development of new methods via rapid prototyping. Most importantly, they are a necessary step towards their use in applications. While considerable advances in the theoretical understanding of probabilistic numerical methods have been made, the lack of high-quality implementations is holding back their adoption. In response, we recently started a community effort to develop an open-source framework named `ProbNum` (http://probnum.org).

A central theme of Dagstuhl Seminars is the open, collaborative atmosphere with a focus on new ideas and tangible outcomes as opposed to existing work. The seminar stimulated multiple focussed discussions around software and additions to ProbNum. Examples included how to best include automatic differentiation functionality, or how to expand the package's Bayesian quadrature functionalities. It also set the starting

point for potential new research collaborations on probabilistic linear solvers and probabilistic numerical methods for PDEs. Even at this point, shortly after the seminar's conclusion, two tangible products are already available: the seminar's participants jointly created a Probabilistic Numerics Wikipedia page https://en.wikipedia.org/wiki/Probabilistic_numerics, and the implementation sessions culminated with a preprint for the community library `ProbNum` [1].

As the organizers we want to thank all participants, both physical and virtual, for their interesting talks, the stimulating discussions and the collaborative overall atmosphere. We also want to thank Schloss Dagstuhl for their technical support that made the challenging hybrid format possible.

References

1 Jonathan Wenger, Nicholas Krämer, Marvin Pförtner, Jonathan Schmidt, Nathanael Bosch, Nina Effenberger, Johannes Zenn, Toni Karvonen Alexandra Gessner, François-Xavier Briol, Maren Mahsereci, and Philipp Hennig. ProbNum: Probabilistic numerics in Python. arxiv preprint, 2021. URL http://arxiv.org/abs/2112.02100

6.38 Adaptive Resource Management for HPC Systems

Organizers: Michael Gerndt, Masaaki Kondo, Barton P. Miller, and Tapasya Patki

Seminar No. 21441

Date: November 1–5, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.1

Participants: Eishi Arima, Eduardo César, Isaías Alberto Comprés Ureña, Michael Gerndt, Jophin John, Matthias Maiterth, Barton P. Miller, Bernd Mohr, Frank Mueller, Santiago Narvaez Rivas, Mirko Rahn, Lubomir Riha, Martin Schulz, Anna Sikora, Ondrej Vysocky, Felix Wolf

Remote Participants: Dong Ahn, Andrea Bartolini, Pete Beckman, Mohak Chadha, Julita Corbalan, Balazs Gerofi, Toshihiro Hanawa, Shantenu Jha, Rashawn Knapp, Masaaki Kondo, Daniel John Milroy, Tapasya Patki, Barry L. Rountree, Roxana Rusitoru, Sakamoto Ryuichi, Wolfgang Schröder-Preikschat

Today's supercomputers have very static resource management. Jobs are submitted via batch scripts to the resource manager, then scheduled on the machine with a fixed set of nodes. Other resources, such as power, network bandwidth and storage are not actively managed and are provided only on a best-effort basis. This inflexible, node-focused and static resource management will have to change in the future due to many reasons, some of them listed below.

First, applications are becoming increasingly more dynamic. Techniques such as adaptive mesh refinement, e.g., as used in Tsunami simulations, lead to scalability changes over the application's execution. Furthermore, only some application phases might profit from specialized accelerators, and I/O phases might even run best with a limited number of compute resources.

Additionally, the execution environment of applications is also becoming dynamic. Modern processors change the clock frequency according to the instruction mix as well as power and thermal envelopes. Heavy use of the vector units can lead to a lower clock frequency to stay in the thermal power budget, for example.

As an independent concern, due to the sheer number of components, failure rates are expected to increase thus slowing down computation or even leading to an increased number of node failures.

Finally, the upcoming machines will be power constrained, which means that the power will have to be carefully distributed among all running applications. The resulting power capping will impact the application's performance due to adaptation of the clock frequency and due to manufacturing variability. These challenges in HPC will only be solvable by using a more adaptive resource management approach. For example, compute nodes need to be redistributed among running applications to adapt to changes in the application's resource requirements either due to a varying number of grid points or interspersed algorithmic phases that profit from certain accelerators; network and I/O bandwidth will have to be assigned to applications to avoid interference caused by contention of concurrent communication and I/O phases; power needs to be dynamically redistributed both within an application and across applications to enable increased efficiency. Dynamic redistribution of resources will also give more flexibility to the resource manager to schedule jobs on the available resources and thus reduce idle times and efficiency lowering contention scenarios, e.g., in the situation of big jobs waiting for execution.

This Dagstuhl Seminar investigated a holistic, layered approach for adaptive resource management. It started with the resource management layer being responsible for scheduling applications on the machine and dynamically allocating resources to the running applications. At the programming level, applications need to be programmed in a resource-aware style such that they can adapt to resource changes and can make most efficient usage of the resources. On top of the programming interfaces, programming tools have to be available that allow the application developers to analyze and tune the applications for the varying amount of available resources. At the application level, applications have to be redesigned to enable significant

gains in efficiency and throughput, e.g., adaptive mesh refinement, approximate computing, and power-aware algorithms are a few aspects to mention here.

The discussions led to a joint summary presenting the state-of-the-art, required techniques on these layers of HPC systems, as well as the foreseen advantages of adaptive resource management.

6.39 Ensuring the Reliability and Robustness of Database Management Systems

Organizers: Alexander Böhm, Maria Christakis, Eric Lo, and Manuel Rigger

Seminar No. 21442

Date: November 1–4, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.20

Participants: Alexander Böhm, Cristian Cadar, Alastair F. Donaldson, Stefania Dumbrava, Marco Guarnieri, Marcel Kost, Burcu Kulahcioglu Ozkan, Hannes Mühleisen, Danica Porobic, Mark Raasveldt, Manuel Rigger, Anupam Sanghi

Remote Participants: Artur Andrzejak, Chee-Yong Chan, Yongheng Chen, Maria Christakis, Jack Clark, Jens Dittrich, Paolo Guagliardo, Jayant R. Haritsa, Miryung Kim, Kyle Kingsbury, Greg Law, Si Liu, Eric Lo, Muhammad Numair Mansur, Zhou Qiang, Tilmann Rabl, Abhik Roychoudhury, Zhendong Su, S. Sudarshan, Tao Xie, Tianyin Xu, Mai Zheng

DataBase Management Systems (DBMSs) are used ubiquitously. Due to the ever-growing number and size of data sets, increasing performance demands, and the virtually unlimited hardware resources that are provided by public cloud infrastructure, sophisticated systems and optimizations are developed continuously. This dynamic and demanding environment is a major challenge for developers of DBMSs, which have to ensure that their systems are both correct and efficient.

Database management systems are a well-established field with several decades of research and engineering attention. These efforts have resulted in a multitude of both open-source and commercial systems that are widely deployed in production today and provide the backbone of a vast range of mission-critical applications. Still, surprisingly, recent work on automatic testing of DBMSs found a large number of bugs in widely-used DBMSs. This clearly indicated that the topic of ensuring the reliability and robustness of DBMS deserves more attention, and that key insights from neighboring domains such as automatic testing and formal methods could potentially help to advance the state of the art in DBMS engineering.

Goals and Outcomes

One of the central goals and outcomes of the seminar was to build a common foundation and understanding for the key challenges of DBMS engineering, and how they can be potentially addressed. To this end, the seminar focused on

- Best practices and challenges in building open source and commercial database engines. Here, the key objectives include a high developer efficiency, mandating quick feedback by tests and verification tools already during feature development, as well as systematic (stress) testing of the software under high load and error conditions.
- The applicability of formal methods and verification tools to DBMS.
  Formal methods can be of great help to prove the correctness of key database system components such as query compilers, distributed consensus protocols, data replication components, or modules dealing with high availability. Still, an important question is how to systematically identify those components that can benefit from formal verification with reasonable implementation effort, and how to best integrate these methods into existing systems.
- Advanced testing techniques such as fuzzers, query synthesis, and workload generators.
  These methods allow to significantly increase the test coverage of a DBMS by systematically exploring uncovered code paths and putting stress on individual, important subsystems such as input verification and error handling that are a frequent source of software defects.
- Methods for the automatic generation of test data and testcase reduction.

Occasionally, defects in database software are only found by customers running very complex queries operating on confidential data sets. Thus, to allow for problem reproduction, developers benefit from a minimal data set and a simplified query specification that does not disclose confidential data or exhibit unnecessary complexity.
- Security aspects such as ensuring confidentiality and data integrity in the presence of different classes of attackers.

Attendee Mix and Seminar Structure

The seminar lasted 2.5 days. Its format and attendee mix was significantly influenced by the ongoing pandemic. Of the 34 attendees, 13 attended in person and 21 remotely. All but one of the in-person attendees were based in Europe. Overall, we received the highest response rate from Europe (20 attendees), and a lower one from Asia (8 attendees) and the US (6 attendees). We are grateful to the two Video Conference Assistants (VCAs), Jack Clark and Mark Raasveldt, who managed the equipment to ensure a smooth experience for all attendees.

We started the seminar with an introduction round in which every attendee introduced themselves. We held another such session in the late afternoon, to accommodate the US attendees. Prior to the seminar, we contacted attendees to give overview talks to establish a common discussion basis, which was useful given that the attendees came from different scientific communities. We had such overview talks on the first and second day. On the second and third day, we had in-depth talks. While we had planned breakout sessions, many of the talks were followed by fruitful and unplanned discussions. On the last day, we had a group discussion on the takeaways and future plans.

Future Plans

One major result from the seminar was to identify open problems and areas of future work that the group wants to address in an interdisciplinary manner. Among others, this includes the creation of a reference manual for database engineering groups to avoid redundant work and re-inventing techniques already established (or discarded) by other teams, the identification of database modules (e.g. the query compiler and transaction processing system) that can benefit from formal verification, designing new test oracles to test various data-centric systems for different kind of bugs, as well as the establishment of a common testcase specification format and a test corpus that can be shared between DBMS engineering teams. We discussed proposing another instance of the Dagstuhl Seminar to utilize the established discussion basis and work on addressing these specific challenges.

6.40 Managing Industrial Control Systems Security Risks for Cyber Insurance

Organizers: Simon Dejung, Mingyan Liu, Arndt Lüder, and Edgar Weippl

Seminar No. 21451

Date: November 7–12, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.36

Participants: Luca Allodi, Gergely Biczók, Rainer Böhme, Carl Denis, Matthias Eckhart, Alexander Horch, Sejdefa Ibisevic, Julia Kittel, Klaus Kursawe, Arndt Lüder, Fabio Massacci, Thomas Steinhaus, Edgar Weippl, Jens Wiesner, Daniel Woods

Remote Participants: Fabrizio Baiardi, Vivien Bilquez, Achim D. Brucker, Richard Clayton, Simon Dejung, Andreas Ekelhart, Barbara Fila, Peter Hacker, Teodore Iazikoff, Helge Janicke, Ersin Kaplan, Marina Krotofil, Éireann Leverett, Mingyan Liu, Markus Maier, Jürgen Musil, Simin Nadjm-Tehrani, Ranjan Pal, Keyun Ruan, Galina Schwartz, Sara Tajik, Josephine Wolff, Gordon Woo, Quanyan Zhu

The security economics community has an ambivalent position on quantifying cyber risk: on the one hand, there is a long-standing interest in establishing cyber risk measurement, while on the other hand, relevant publications report contradictory results [7] or present models with insufficient evidence of validity [6]. As researchers remain cautious about the interpretation of modeling results, the insurance industry's need for gaining a quantitative understanding of cyber risk is more critical than ever. A serious concern for insurers is that a large-scale cyber-attack could result in significant claims arising from transferred security risks. Considering the adverse business implications of cyber-related high-impact, low-frequency events, it is necessary to estimate and control the potential exposure to such losses. Previous industry studies attempted to address this issue by proposing hypothetical worst-case scenarios involving power grids [8], ports [9], and other industrial applications [4]. Furthermore, several industry-led working groups conducted workshops to assess the plausibility, impact, and claim implications of potential ICS-related catastrophic loss events (cf., for instance, [5, 10–12]). However, it is still not fully understood how the peculiarities of ICSs, technological change in light of strategic initiatives (e.g., Industry 4.0 [3]), and the increasingly sophisticated nature of cyber-physical attacks influence the loss frequency and loss severity [1]. Moreover, a holistic consideration of cyber-physical risk featuring the complete ICS lifecycle calls for an interdisciplinary research approach [2].

Thus, the aim of this Dagstuhl Seminar was to bring together different communities to foster research activities that advance the understanding of cyber risks pertaining to ICSs and associated insurance aspects. The concepts developed as part of this seminar are a result of interdisciplinary work conducted by academics and industry professionals, both junior and senior, from the fields of

(i) computer science,
(ii) automation engineering,
(iii) actuarial science, and
(iv) economics.

To address the issues outlined above, the purpose of the seminar was to make the first steps toward a probabilistic cyber catastrophe model that is tailored to the ICS domain. In particular, we planned to achieve the basis of an economic loss model that builds upon worst-case scenarios in which globally and simultaneously many industrial processes in critical infrastructure sectors (e.g., power, petrochemical, transport, logistics) are affected by cyber-attacks. Figure 6.13 visualizes possible components of such a model, which were discussed and challenged during the seminar. In the first phase, the scenario is formulated, fundamental assumptions are specified, and the theoretical basis of the statistical model is formed. After that, the model is calibrated with data obtained from various sources, such as loss databases or subject-matter experts. Finally, the model and its underlying assumptions are validated.

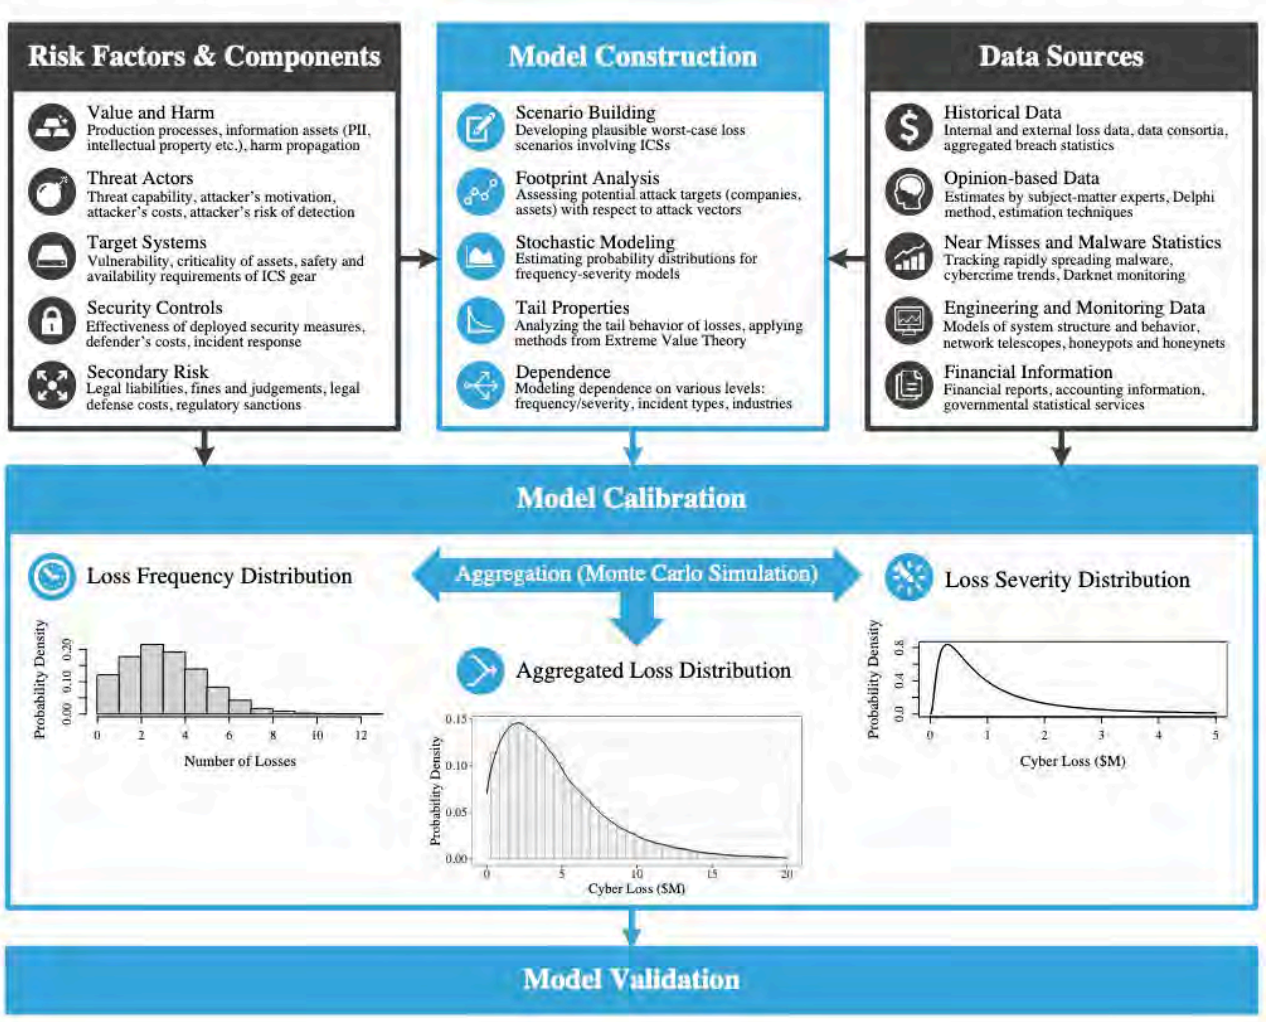

Fig. 6.13
Potential components of a probabilistic cyber catastrophe model for the ICS domain (adapted from materials provided by SCOR SE).

To set the frame for the seminar, the organizers defined four topics that were covered in plenary sessions and breakout sessions. In each plenary session, lightning talks were held that motivated the collaborative work in the breakout sessions. The working groups studied the same overarching topic of the breakout session (yet each with a different focus) in order to strengthen interdisciplinary exchange.

Overall, the following topics and motivating research questions were addressed:

1. *ICS Threat Landscape*: How have cyber attacks against ICSs evolved and what should we expect in terms of attack sophistication, persistence, and impact in the future?
2. *Cyber-Physical Risk Quantification*: How can we quantitatively model economic losses caused by ICS-focused cyber risks (i.e., probabilistic cyber catastrophe model)?
3. *Insurance*: What are the opportunities and limitations of transferring ICS-focused cyber risks to insurers?
4. *Management of Security Risks*: Which hard (e.g., technological security measures) and soft (e.g., information sharing, regulations, funding) factors increase or reduce the attack likelihood and severity?

The seminar started with a welcome session to bridge the disciplinary gap. In this session, the organizers presented the seminar program, explained key terms, and discussed core concepts to familiarize attendees with the terminologies used by different communities. Over the following days, several participants gave lightning talks that focused on the following topics:

- cyber-physical systems, security-relevant aspects within their lifecycle, and procurement considerations,
- current ICS security challenges with an emphasis on technological trends (e.g., Industry 4.0, smart manufacturing, Industrial Internet of Things),
- cyber-physical risk assessments, where special attention was given to analysis and quantification methods,
- various aspects of (cyber) insurance (e.g., cyber cat modeling, underwriting, economic problems, regulations), and
- security economics, featuring studies on cybercrime analysis and vulnerability forecasting.

The lightning talks gave participants the opportunity to present new perspectives and challenges, which led to lively discussions that shaped the group sessions. Unfortunately, the restrictions caused by the SARS-CoV-2 pandemic made it not possible to conduct the estimation exercises with all participants, which would have been required for achieving the cyber cat model. However, conducting the seminar in a hybrid format still enabled the participants both on-site and remote to work together on challenging open questions, contribute to group discussions, and forge new research collaborations.

The organizers thank all participants for their valuable contribution. Furthermore, this seminar would not have been possible without the great technical support provided by the Schloss Dagstuhl staff and the considerable effort made by the video conferencing assistants Sejdefa Ibisevic, Markus Maier, and Sara Tajik.

References

1 Matthias Eckhart, Bernhard Brenner, Andreas Ekelhart, and Edgar Weippl. Quantitative security risk assessment for industrial control systems: Research opportunities and challenges. *Journal of Internet Services and Information Security (JISIS)*, 9(3):52–73, August 2019.

2 Gregory Falco, Martin Eling, Danielle Jablanski, Matthias Weber, Virginia Miller, Lawrence A. Gordon, Shaun Shuxun Wang, Joan Schmit, Russell Thomas, Mauro Elvedi, Thomas Maillart, Emy Donavan, Simon Dejung, Eric Durand, Franklin Nutter, Uzi Scheffer, Gil Arazi, Gilbert Ohana, and Herbert Lin. Cyber risk research impeded by disciplinary barriers. *Science*, 366(6469):1066–1069, 2019.

3 Henning Kagermann, Johannes Helbig, Ariane Hellinger, and Wolfgang Wahlster. Recommendations for implementing the strategic initiative INDUSTRIE 4.0 – securing the future of german manufacturing industry. Final report of the Industrie 4.0 working group, acatech – National Academy of Science and Engineering, München, April 2013.

4 Lloyd's of London, Guy Carpenter, and CyberCube Analytics. Cyber risk: The emerging cyber threat to industrial control systems. Technical report, Lloyd's of London, Guy Carpenter, and CyberCube Analytics, February 2021.

5 Lobo, Francis. Upstream oil & gas cyber risk: Insurance technical review. Technical report, Joint Rig Committee, May 2018. A Joint Rig Committee Report.

6 Vilhelm Verendel. Quantified security is a weak hypothesis: A critical survey of results and assumptions. In *Proceedings of the 2009 Workshop on New Security Paradigms Workshop*, NSPW '09, pages 37–50, New York, NY, USA, 2009. ACM.

7 Daniel W. Woods and Rainer Böhme. SoK: Quantifying cyber risk. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 211–228, May 2021.

8 Lloyd's of London and Cambridge Centre for Risk Studies. Business blackout: The insurance implications of a cyber attack on the us power grid. Technical report, Lloyd's of London and Cambridge Centre for Risk Studies, July 2015.

9 Lloyd's of London, Cambridge Centre for Risk Studies, and Nanyang Technological University. Shen attack: Cyber risk in asia pacific ports. Technical report, Lloyd's of London, Cambridge Centre for Risk Studies, and Nanyang Technological University, 2019.

10 Dejung, Simon. Economic impact of cyber accumulation scenarios. Technical report, SCOR Global P&C, 2017.

11 Dejung, Simon. Newsletter – risk assessment for ICS/SCADA security in industrial property, engineering, power, oil & gas. Technical report, SCOR Global P&C, March 2018. A joint workshop in March 2018 by LMA, IMIA & OPERA at SCOR (Zurich).

12 IMIA Working Group. Cyber risks: Engineering insurers perspective. Technical Report 98 (16), September 2016. IMIA Annual Conference 2016 – Doha, Qatar.

6.41 Unambiguity in Automata Theory

Organizers: Thomas Colcombet, Karin Quaas, and Michał Skrzypczak

Seminar No. 21452

Date: November 7–12, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.57

Participants: Achim Blumensath, Udi Boker, Dmitry Chistikov, Lorenzo Clemente, Thomas Colcombet, Wojciech Czerwinski, Diego Figueira, Emmanuel Filiot, Simon Jantsch, Ismaël Jecker, Stefan Kiefer, Shankaranarayanan Krishna, Denis Kuperberg, Karoliina Lehtinen, Antoine Mottet, Anca Muscholl, Pierre Ohlmann, Guillermo A. Pérez, Jakob Piribauer, Gabriele Puppis, Karin Quaas, Alexander Rabinovich, Michael Raskin, Mahsa Shirmohammadi, Michal Skrzypczak, Sarah Winter, Georg Zetzsche

Remote Participants: Christel Baier, Johanna Björklund, Michaël Cadilhac, Antonio Casares, Nathanael Fijalkow, Mika Göös, Arthur Jaquard, Stefan Kiefer, Christof Löding, Sylvain Lombardy, Radek Piórkowski, Mikhail V. Volkov

The Dagstuhl Seminar 21452 "Unambiguity in Automata Theory" was a seminar of five days that took place from November 7th to 12th, 2021, organized by Thomas Colcombet, Karin Quaas, and Michał Skrzypczak. A general goal of the seminar was to bring together experts from different fields of automata theory, to stimulate an exchange of recent results and new proof techniques concerning unambiguity and related topics from automata theory. There were 26 on-site participants from nine different countries (Belgium, Czech Republic, France, Germany, India, Italy, Poland, UK), and further 10 remote participants from seven countries (France, Germany, Poland, Sweden, Switzerland, UK, USA).

The central topic of the seminar was *unambiguous automata*. An automaton is unambiguous if it can make nondeterministic choices, but it is guaranteed that for every input there is *at most one accepting* run. There have recently been numerous new results concerning unambiguous automata; at the same time, a lot of natural and interesting problems have been open for decades. Before the seminar, we identified the following key topics/open problems:

- **Unambiguous Finite Automata** What is the state complexity of the complementation of unambiguous automata? Here, the state complexity refers to how the number of states of the resulting automaton depends on the number of states of the original automaton.
- **Unambiguous versions of infinite state systems, such as vector addition systems with states (VASS) or register automata** Open problems concerning such systems are, for instance: What can be new techniques for proving lower bounds for the containment problem? Are languages accepted by unambiguous register automata with guessing closed under complement?
- **Unambiguous tree automata** One of the most important open questions is how to decide whether a given tree-regular language is recognizable by an unambiguous automaton.
- **Büchi automata and probabilistic automata** What is the computational complexity of the containment of unambiguous Büchi automata?
- **Tropical automata** For this class of weighted automata one of the most important and long standing open questions is whether a given series is polynomially ambiguous.

The seminar was planned to consist of talks and working group sessions, where participants could work on site on open problems. In order to integrate all participants and to initiate new collaborations, we started the seminar on Monday with introductory talks, where every participant shortly introduced herself to the group. In these introductory sessions, it was also possible to announce open problems the participants were interested to work on during the seminar. We had additionally collected such open problems before the seminar to make them available to the participants in advance.

The second day of the seminar (Tuesday) was dedicated to presentations given by the participants. This day started with an invited talk by Denis Kuperberg on good-for-games automata. Later the day, eight participants of the seminar presented short contributed talks on topics related to unambiguity.

Wednesday began with the invited talk by Gabriele Puppis on register automata. Later, a single contributed talk was given and the whole afternoon was devoted to an excursion and group work.

On Thursday morning, Wojtek Czerwiński gave an invited talk on future-determinisation. After that, four contributed talks were given, and the late afternoon was devoted to work in subgroups.

Finally, on Friday morning we held a closing ceremony. The rest of the day was left to participants to summarise their discussions in subgroups and prepare for departure.

During all days, we have used Schloss Dagstuhl's excellent technical facilities to connect and communicate to remote participants of the seminar. Our experiences regarding such a hybrid Dagstuhl Seminar are twofold. On the one hand, it is practical to give remote participants the opportunity to follow the on-site presentations (and Sylvain Lombardi also gave a remote talk). On the other hand, our main aim was to bring together researchers to actually work on concrete problems. It was difficult to integrate participants in group work, when groups gather at different places in the facilities, or when important discussions are led during the excursion or the dinner. We appreciated very much the opportunity to gather on site at Schloss Dagstuhl after a long time of only non-physical meetings due to the Covid pandemics. As summarized in Section 4 of the full report, several new collaborations between participants of the seminar have been initiated. We hope that the seminar has inspired new ideas, and interesting new results will be published by the participants.

We would like to warmly thank Schloss Dagstuhl for making this seminar possible. We especially would like to thank for the great help and support in the organization before and during the seminar.

6.42 Descriptive Set Theory and Computable Topology

Organizers: Mathieu Hoyrup, Arno Pauly, Victor Selivanov, and Mariya I. Soskova

Seminar No. 21461

Date: November 14–19, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.72

Participants: Vasco Brattka, Riccardo Camerlo, Raphael Carroy, Jacques Duparc, Ivan Georgiev, Jun Le Goh, Vassilios Gregoriades, Mathieu Hoyrup, Steffen Lempp, Elvira Mayordomo, Russell G. Miller, Eike Neumann, Adam Ó Conghaile, Arno Pauly, Marcin Sabok, Matthias Schröder, Alexandra A. Soskova, Martin Ziegler

Remote Participants: Alessandro Andretta, Nikolay Bazhenov, Ruiyuan (Ronnie) Chen, Matthew de Brecht, Damir D. Dzhafarov, Olivier Finkel, Ekaterina Fokina, Johanna N. Y. Franklin, Reinhold Heckmann, Daisuke Ikegami, Corrie Ingall, Josiah Jacobsen-Grocott, Iskander Shagitovich Kalimullin, Takayuki Kihara, Margarita Korovina, Davorin Lesnik, Neil Lutz, Alexander Melnikov, Joseph S. Miller, Luca Motto Ros, Takako Nemoto, Keng Meng Ng, André Otfrid Nies, Alexey Ostrovsky, Jan Reimann, Philipp Schlicht, Victor Selivanov, Svetlana Selivanova, Alex Simpson, Theodore A. Slaman, Mariya I. Soskova, Daniel Turetsky, Linda Westrick

Research area and topics

Descriptive set theory traditionally studies the complexity of subsets of and functions between Polish spaces (which are the completely metrizable separable spaces). As a mathematical area, it has well-established interactions with set theory and real analysis. Its canonical textbook is Kechris [11].

Following the developments in (classical) descriptive set theory, also the area of effective descriptive set theory flourished. In a way, this is the result of replacing *continuous* by *computable* everywhere, and by replacing arbitrary countable union by effective ones. Here, the canonical textbook is Moschovakis' [19]. While classical descriptive set theory is trivial on discrete spaces, the results from effective descriptive set theory on $\mathbb{N}$ often generalize results from computability theory. While this is rarely emphasized (see [20] for an exception), one can recover classical descriptive set theory from effective descriptive set theory by relativization – provided that theorems are phrased in the right way.

Recent years have seen a lot of interest in the interplay between descriptive set theory and theoretical computer science going beyond the natural meeting point of effective DST. Four core developments outlined below are particularly relevant for the meeting:

DST on spaces of interest for TCS

Certain classes of topological spaces were revealed as applicable to reasoning about the semantics of programming. The most famous example is domain theory, but Escardo's synthetic topology [6] or the relationship between well-structured transition systems and Noetherian spaces revealed by Goubault-Larrecq [7] were also very influential. The spaces relevant for TCS are often not Hausdorff, and in particular not Polish. Selivanov pioneered the call for a development of descriptive set theory for these spaces [28, 29]. A break-through was achieved by de Brecht [3] with identifying the class of quasi-Polish spaces as a common generalization of Polish spaces and omega-continuous domains, and by showing that many core results of descriptive set theory can be extended to quasi-Polish spaces.

In computable analysis, we typically work with the category of admissible represented spaces (equivalently, with QCB_0-spaces, i.e. T_0-quotients of countably-based spaces) [24]. This is a Cartesian-closed category, meaning that we can form

function spaces. This is a very natural requirement from a TCS-perspective, but does not preserve being countably-based. How descriptive set theory works on non-countably-based spaces is still a mystery. de Brecht, Selivanov and Schröder have undertaken initial investigations, in particular into the Kleene-Kreisel spaces in [5, 26, 27]. Hoyrup has shown that even very simple non-countably-based spaces such as $\mathcal{O}(\omega^\omega)$ exhibit very unfamiliar behaviour compared to the usual DST [8].

Synthetic DST

de Brecht and Pauly observed a connection between synthetic topology (which in turn can be seen as the theory of functional programming [6]), models of hypercomputation and descriptive set theory [4,22,23]. This connection opens up the opportunity to apply reasoning styles about models of computation to descriptive set theory. Work by Kihara on the Jayne-Rogers conjecture has shown significant potential of this approach for solving open questions in descriptive set theory [12]. There is also a hope that this theory can connect to other parts of TCS such as descriptive complexity.

DST and computability theory

Traditional computability theory, in particular the study of enumeration degrees, was related to the study of topological spaces via the notion of point degree spectrum introduced by Kihara and Pauly [13], building on earlier work by J. Miller [17]. This lets us reason about the degrees of individual points in a topological space, and understand properties of the space in terms of what degrees are realized there. This technique was already used to resolve a long-standing open question by Jayne ([9], also [21]) on the number of sigma-homeomorphism types of Polish spaces in [13].

This connection is bidirectional, and also allows for the application of topological arguments in computability theory. As such, it has inspired a flurry of recent developments in the area of enumeration degrees by J. Miller, M. Soskova and others [1, 2, 16, 18]. Particularly remarkable here is the existence of non-total almost-total enumeration degrees. This is a purely recursion-theoretic statement, but the various known proofs all invoke topological arguments such as Brouwer's Fixed Point theorem, Urysohn's metrization theorem or Hurewicz' and Wallmann's characterization of countably-dimensional Polish spaces.

Of a similar flavour (but the precise connections are still unclear) is the approach to fractal geometry and Hausdorff dimension via *effective dimension* of points, defined via Kolmogorov complexity [14]. This approach has already been demonstrated to provide strengthening of core results of fractal geometry, in many cases by rendering inessential restrictions to measurable sets. This includes a reproof of known answer to the two-dimensional Kakeya-conjecture [15].

coPolish spaces and computational complexity

In general, it seems that computational complexity of algorithms from computable analysis needs second-order complexity (Kawamura and Cook [10]). For certain spaces, however, runtimes of algorithms are still first-order objects [25]. Ongoing work by de Brecht and Schröder has shown that this holds for the coPolish spaces, a dual notion to the quasi-Polish spaces. As such, it seems that "spaces where descriptive set theory is well-behaved" is the dual notion to "spaces where complexity theory is well-behaved". This merits further attention by a broader community.

Seminar structure

As our seminar brought together researchers from previously rather disconnected areas, we included several tutorial talks of one hour each to introduce the various facets of our seminar topics to everyone. The talks covered *Quasi-Polish spaces* (Matthew de Brecht), *Quantitative Coding and Complexity Theory of Continuous Data* (Martin Ziegler), *CoPolish spaces and Effectively Hausdorff spaces* (Matthias Schröder), *New directions in Synthetic Descriptive Set Theory* (Takayuki Kihara), *Categorical aspects of Descriptive Set Theory* (Ruiyuan Chen), *Topology reflected in the enumeration degrees* (Joseph S. Miller), *Point-free Descriptive Set Theory* (Alex Simpson) and *Borel combinatorics fail in HYP* (Linda Westrick).

In addition, we had many short (fifteen minute) talks introducing topics or open questions. The prompt for these talks was "What theorem do you want to prove during/following this workshop?", and we are excited to learn what will come from this in the next months.

Challenges in hybrid Dagstuhl meetings

While the organizers and most participants had grown very accustomed to virtual meetings, the setting for our seminar was decidedly hybrid: About half of the participants were present in person, half were participating remotely. The same split applied to the organizing team.

The Dagstuhl team had equipped our main meeting room with multiple cameras and microphones (including microphones suspended from the ceiling throughout the room to pick up audience contributions). The equipment was controlled by several volunteers amongst the participants, and we are very grateful to Nikolay Bazhenov, Josiah Jacobsen-Grocott and Eike Neumann for having performed this crucial role. This setup made interactions in the lecture theatre between remote and in-person participants almost seamless.

A feature we felt was both crucial for a successful Dagstuhl Seminar and difficult to accomplish in a hybrid setting are the informal discussions taking place in smaller groups. Our approach was to make those slightly less informal, and to use the collaboration platform Slack for arranging meetings. Slack also served for asking questions somewhat after the talks. This was somewhat successful, and several fruitful discussions involving both remote and in-person participants took place. It is difficult to ascertain though how much potential for additional discussions remained untapped.

References

1 Uri Andrews, Hristo A. Ganchev, Rutger Kuyper, Steffen Lempp, Joseph S. Miller, Alexandra A. Soskova, and Mariya I. Soskova. On cototality and the skip operator in the enumeration degrees. preprint.

2 Uri Andrews, Gregory Igusa, Joseph S. Miller, and Mariya I. Soskova. Characterizing the continuous degrees. *Israel Journal of Mathematics*, 234:743–767, 2019.

3 Matthew de Brecht. Quasi-Polish spaces. *Annals of Pure and Applied Logic*, 164(3):354–381, 2013.

4 Matthew de Brecht and Arno Pauly. Noetherian Quasi-Polish spaces. In Valentin Goranko and Mads Dam, editors, *26th EACSL Annual Conference on Computer Science Logic (CSL 2017)*, volume 82 of *LIPIcs*, pages 16:1–16:17. Schloss Dagstuhl, 2017.

5 Matthew de Brecht, Matthias Schröder, and Victor Selivanov. Base-complexity classifications of QCB_0-spaces. In Arnold Beckmann, Victor Mitrana, and Mariya Soskova, editors, *Evolving Computability*, pages 156–166. Springer, 2015.

6 Martín Escardó. Synthetic topology of datatypes and classical spaces. *Electronic Notes in Theoretical Computer Science*, 87, 2004.

7 Jean Goubault-Larrecq. *Non-Hausdorff Topology and Domain Theory*. New Mathematical Monographs. Cambridge University Press, 2013.

8 Mathieu Hoyrup. Results in descriptive set theory on some represented spaces. arXiv 1712.03680, 2017.

9 J. E. Jayne. The space of class α Baire functions. *Bull. Amer. Math. Soc.*, 80:1151–1156, 1974.

10 Akitoshi Kawamura and Stephen Cook. Complexity theory for operators in analysis. *ACM Transactions on Computation Theory*, 4(2), 2012.

11 A.S. Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer, 1995.

12 Takayuki Kihara. Decomposing Borel functions using the Shore-Slaman join theorem. *Fundamenta Mathematicae*, 230, 2015. arXiv 1304.0698.

13 Takayuki Kihara and Arno Pauly. Point degree spectra of represented spaces. arXiv:1405.6866, 2014.

14 Jack Lutz. The dimensions of individual strings and sequences. *Information and Computation*, 187:49–79, 2003.

15 Jack H. Lutz and Neil Lutz. Algorithmic Information, Plane Kakeya Sets, and Conditional Dimension. In Heribert Vollmer and Brigitte Valle´e, editors, *34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)*, volume 66 of *LIPIcs*, pages 53:1–53:13. Schloss Dagstuhl, 2017.

16 Ethan McCarthy. Cototal enumeration degrees and their application to computable mathematics. *Proceedings of the AMS*, 146:3541–3552, 2018.

17 Joseph S. Miller. Degrees of unsolvability of continuous functions. *Journal of Symbolic Logic*, 69(2):555 – 584, 2004.

18 Joseph S. Miller and Mariya I. Soskova. Density of the cototal enumeration degrees. *Annals of Pure and Applied Logic*, 2018.

19 Yiannis N. Moschovakis. *Descriptive Set Theory*, volume 100 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, 1980.

20 Yiannis N. Moschovakis. Classical descriptive set theory as a refinement of effective descriptive set theory. *Annals of Pure and Applied Logic*, 162:243–255, 2010.

21 Luca Motto Ros, Philipp Schlicht, and Victor Selivanov. Wadge-like reducibilities on arbitrary quasi-polish spaces. *Mathematical Structures in Computer Science*, pages 1–50, 11 2014. arXiv 1204.5338.

22 Arno Pauly and Matthew de Brecht. Non-deterministic computation and the Jayne Rogers theorem. *Electronic Proceedings in Theoretical Computer Science*, 143, 2014. DCM 2012.

23 Arno Pauly and Matthew de Brecht. Descriptive set theory in the category of represented spaces. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 438–449, 2015.

24 Matthias Schröder. Extended admissibility. *Theoretical Computer Science*, 284(2):519–538, 2002.

25 Matthias Schröder. Spaces allowing type-2 complexity theory revisited. *Mathematical Logic Quarterly*, 50(4/5):443–459, 2004.

26 Matthias Schröder and Victor Selivanov. Hyperprojective hierarchy of QCB_0-spaces. *Computability*, 4, 2015. arXiv 1404.0297.

27 Matthias Schröder and Victor L. Selivanov. Some hierarchies of QCB_0-spaces. *Mathematical Structures in Computer Science*, 25(8):1799–1823, 2015. arXiv 1304.1647.

28 Victor L. Selivanov. Difference hierarchy in φ-spaces. *Algebra and Logic*, 43(4):238–248, 2004.

29 Victor L. Selivanov. Towards a descriptive set theory for domain-like structures. *Theoretical Computer Science*, 365(3):258–282, 2006.

6.43 Foundations of Persistent Programming

Organizers: Hans-J. Boehm, Ori Lahav, and Azalea Raad

Seminar No. 21462

Date: November 14–17, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.94

Participants: Mark Batty, Hans-J. Boehm, Michael D. Bond, Brijesh Dongol, Michal Friedman, Michalis Kokologiannakis, Ori Lahav, Maged M. Michael, Erez Petrank, Azalea Raad, Konstantinos Sagonas, Michael Scott, Viktor Vafeiadis, William Wang, Heike Wehrheim

Remote Participants: Hagit Attiya, Parosh Aziz Abdulla, Piotr Balcer, Eleni Bila, Dhruva Chakrabarti, Soham Chakraborty, Ricardo Ciríaco da Graca, Brian Demsky, Derek Dreyer, João F. Ferreira, Maurice Herlihy, Joseph Izraelevitz, Limin Jia, Jeehoon Kang, Artem Khyzha, Umang Mathur, Paul McKenney, Nikos Nikoleris, Andy Rudoff, Peter Sewell, John Wickerson

We brought together 15 in-person attendees at Schloss Dagstuhl, with a roughly equal number of remote attendees. Remote attendance was challenging, particularly for attendees from very different time-zones. It nonetheless provided the opportunity for us to hear from a wider selection of participants.

We decided up-front not to try to cater the schedule to remote participation. Given the number of time zones covered by the participants, we continue to believe that, although it clearly had adverse impacts, it was the right decision.

We had a number of remote presentations that included interactions with the speakers. Otherwise the discussion tended to happen mostly among the in-person participants, in spite of the excellent AV systems at Dagstuhl. Many remote participants were limited in attendance due to time-zone issues.

Our area is perhaps unique, in that it includes deep theoretical work, but is also very dependent on technological developments. Accordingly, the participants of the seminar were from a spectrum of topics ranging from theory of distributed systems to hardware specification and design. The seminar gave many of us the opportunity to catch up on both theory and practice, including input from some participants with more direct insights into industrial developments.

We began the seminar by reviewing some of the underlying assumptions that were made by prior work in this field, often without certainty about their correctness. We were actually able to get much more shared clarity on a few of these as a result of audience discussion during the seminar. Some of us learned that non-volatile caches are being publicly discussed by Intel, and that there also is similar agreement that memory encryption, to restore volatility when needed, is desirable. We also learned that writes to the same cache line are not just believed by software researchers to reach memory in the correct order, but at least one hardware vendor also agrees. Though this last fact is rather obscure, it is important for some NVM algorithms, and not normally reflected in hardware manuals.

The rest of the seminar consisted of talks and group discussions. Three talks were longer overview talks on different aspects (Michael Scott on buffered persistency, Parosh Aziz Abdulla on verification, and Erez Petrank on persistent lock-free data structures). We did not feel the need for smaller break-out sessions, since the in-person group was quite small, largely due to our timing with respect to Covid waves. Much of the benefit here appears to have been in listening to discussions that often were either significantly more theoretical or significantly more practical than our own research.

NVM programming is both complicated by, and often synergistic with concurrent programming. Much of our focus was on the interaction between the two. Due to these close interactions, we asked several speakers to talk about concurrency issues that seemed particularly relevant (e.g., Peter Sewell on Armv8-A virtual memory model, Mark Batty on novel solution to the "thin air problem", and Paul McKenney on weak memory schemes used in the Linux kernel).

Several talks raised questions on the foundations of the field,

such as what are the hardware-supplied programming models, or what it means for a persistent program to be correct. It is entirely possible that most future NVM programmers will be more concerned with something like the persistent transactions discussed in Michael Bond's talk. However, given the relative immaturity of the area and foundational uncertainty, the emphasis seemed appropriate.

6.44 Geometric Modeling: Interoperability and New Challenges

Organizers: Falai Chen, Tor Dokken, and Géraldine Morin

Seminar No. 21471

Date: November 21–26, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.111

Participants: Arturs Berzins, Tor Dokken, Konstantinos Gavriil, Thomas A. Grandine, Stefanie Hahmann, Rene Hiemstra, Bert Jüttler, Panagiotis Kaklis, Rimvydas Krasauskas, Sylvain Lefebvre, Tom Lyche, Carla Manni, Géraldine Morin, Helmut Pottmann, Ulrich Reif, Espen Sande, Hendrik Speleers, Deepesh Toshniwal

Remote Participants: Gudrun Albrecht, Falai Chen, Ilke Demir, Nira Dyn, Gershon Elber, Carlotta Giannelli, Ron Goldman, Hans Hagen, Qi-xing Huang, Xiaohong Jia, H (Alicia) Kim, Tae-wan Kim, Jiri Kosinka, Yang Liu, Nicolas Mellado, Suraj R. Musuvathy, Francesco Patrizi, Jörg Peters, Konrad Polthier, Jeff Poskin, Xiaoping Qian, Ernst Rank, Elissa Ross, Malcolm A. Sabin, Péter Salvi, Maria Lucia Sampoli, Giancarlo Sangalli, Hiromasa Suzuki, Tamas Várady, Nelly Villamizar, Wenping Wang, Juyong Zhang, Yongjie Jessica Zhang

The Dagstuhl Seminar, initially planned in May 2020, took place as a hybrid conference in November 2021. Eighteen participants were on site, and thirty three participated remotely out of which five from East Asia and twelve from America.

Due to the pandemic, getting together for a conference has been an important event, and an outstanding exchange time between researchers (compared to a two years pandemic context where interaction has greatly been reduced). In particular, having a significant part of the participants on site has been a real asset compared to the full online conferences. Note also that this has been particularly true for young researchers that are in the process of developing networks and developing collaborations.

48 talks were given including 18 on site and 30 remotely. The program was organized into topics and structured to the extent possible to minimize the challenges posed by the time difference between on-site and remote participants. Speakers from East-Asia were assigned time slots in the morning and speakers from America in the afternoon. The beginning of the afternoon was a privileged time for all participants to meet. The social afternoon was canceled, as it would not have been inclusive for remote participants.

The time freed allowed us to extend the time assigned to topic focused groups sessions. This triggered, under the supervision of five on-site participants, development of topic focused reports. Two of these reports, *The Future of CAD* (group led by Tom Grandine) and *Design Optimization* (group led by Konstantinos Gavriil) address the evolution of the application fields in Geometric Modeling, closely linked to its use in Industry. Three other reports on emerging topics have also been based on the group working sessions. *Additive Manufacturing* (a group led by Sylvain Lefebvre) has been identified as a disruptive technology and has triggered the emergence of new geometric models and materials. *Isogeometric Analysis* (group led by Carla Manni) addresses how the gap between geometric modeling and simulation can be bridged by replacing the traditional shape functions of Finite Element Analysis by B-splines that cross element boundaries. It thus supplies continuous models connecting the representations of Computer Aided Design and Finite Element Analysis. *Geometric Machine Learning* (group led by Rene Hiemstra) is a fast evolving domain. Deep learning approaches have already changed the field of Computer Vision, and the contribution into Geometric Modeling is becoming more pregnant. These reports offer to the participants, and beyond, a perspective of the coming challenges in the field of Geometric Modeling.

On top of the communications done in Dagstuhl, a special issue of the journal Graphical Models, has been planned. Submission to the journal is pending.

6.45 Geometric Logic, Constructivisation, and Automated Theorem Proving

Organizers: Thierry Coquand, Hajime Ishihara, Sara Negri, and Peter M. Schuster

Seminar No. 21472

Date: November 21–26, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.151

Participants: Karim Johannes Becher, Arnold Beckmann, Ulrich Berger, Ulrik Buchholtz, Gabriele Buriola, Giulio Fellin, Anton Freund, Matthias Hutzler, Rosalie Iemhoff, Ulrich Kohlenbach, Henri Lombardi, Julien Narboux, Stefan Neuwirth, Eugenio Orlandelli, Iosif Petrakis, Morgan Rogers, Peter M. Schuster, Matteo Tesi

Remote Participants: Jan Belle, Marc Bezem, Pierre Boutry, Olivia Caramello, Liron Cohen, Thierry Coquand, Laura Crosilla, Tiziano Dalmonte, Makoto Fujiwara, Hajime Ishihara, Predrag Janicic, Tatsuji Kawai, Vesna Marinkovic, Kenju Miyamato, Sara Negri, Takako Nemoto, Satoru Niki, Edi Pavlovic, Cosimo Perini Brogi, Thomas Powell, Michael Rathjen, Helmut Schwichtenberg, Monika Seisenberger, Sana Stojanovic-Djurdjevic, Daniel Wessel, Chuangjie Xu

A central question has remained from the foundational crisis of mathematics about a century ago: What is the extent of intuitionistic mathematics? From a practical angle, the question is to which extent any given proof is effective, which proofs of which theorems can be rendered effective, and whether and how numerical information such as bounds and algorithms can be extracted from proofs. Ideally, all this is treated by manipulating proofs mechanically and/or by adequate proof-theoretic metatheorems (proof translations, automated theorem proving, program extraction from proofs, proof analysis, proof mining, etc.). In this vein, the central question should rather be put as follows: What is the computational content of proofs?

Guided by this form of the central question, the Dagstuhl Seminar 21472 put a special focus on coherent and geometric theories and their generalisations. These indeed are fairly widespread in mathematics and non-classical logics such as temporal and modal logics, a priori amenable for constructivisation in the vein of Barr's Theorem, and particularly suited as a basis for automated theorem proving. Specific topics included categorical semantics for geometric theories, complexity issues of and algorithms for geometrisation of theories with the related speed-up questions, the use of geometric theories in constructive mathematics up to finding algorithms, proof-theoretic presentation of sheaf models and higher toposes, and coherent logic for automated proving.

The Dagstuhl Seminar 21472 attracted researchers and practitioners from all over the world, including participants from various research areas in order to broaden the scope of the seminar and to create connections between communities. The seminar participants presented and discussed their research by means of programmed and ad-hoc talks, and a tutorial on Agda the well developed proof assistant based on dependent type theory – was held over several time slots. Numerous new research directions were developed in small working groups: for example, new perspectives on classifying toposes in algebraic geometry, applications of dynamical methods to quadratic forms, and Zorn induction to capture transfinite methods computationally.

The tireless efforts by Dagstuhl staff notwithstanding, it would not be fair to say that this seminar did not suffer from the pandemic-related travel restrictions by which many invitees were confined to remote participation, which of course made hard if not impossible that they took part at the invaluable informal exchange on-site characteristic for events held at Dagstuhl. Under the given circumstances, however, the seminar was still judged a success by all the participants. Following an unconditional request by many, the organisers intend to propose a follow-up Dagstuhl Seminar on a related topic in the near future – if possible, all on-site.

6.46 Secure Compilation

Organizers: David Chisnall, Deepak Garg, Catalin Hritcu, and Mathias Payer

Seminar No. 21481

Date: November 28 to December 3, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.10.173

Participants: Roberto Blanco, Stefan Brunthaler, Matteo Busi, Dominique Devriese, Akram El-Korashy, Deepak Garg, Anitha Gollamudi, Marco Guarnieri, Catalin Hritcu, Marco Patrignani, Jan Reineke, Shweta Shinde, Jeremy Thibault, Thomas Van Strydonck, Ingrid Verbauwhede

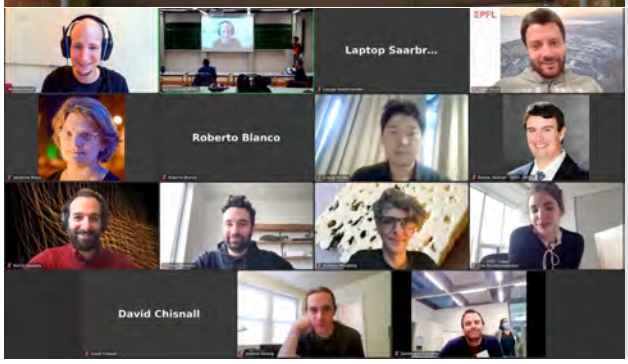

Remote Participants: Amal Ahmed, Arthur Azevedo de Amorim, Gilles Barthe, Joseph Bialek, Sandrine Blazy, Nathan Burow, David Chisnall, Mads Dam, Ergys Dona, Cédric Fournet, Tal Garfinkel, Chung-Kil Hur, Jérémie Koenig, Per Larsen, Amit Levy, Toby Murray, Andrew Myers, Santosh Nagarakatte, Elisabeth Oswald, Zoe Paraskevopoulou, Mathias Payer, Andreas Rossberg, Kostya Serebryany, Peter Sewell, Zhong Shao, Deian Stefan, Gang Tan, Nikos Vasilakis, Marco Vassena, Drew Zagieboylo

Secure compilation is an emerging field that puts together advances in security, programming languages, compilers, systems, verification, and hardware architectures to devise compilation chains that eliminate security vulnerabilities, and allow sound reasoning about security properties in the source language. For example, all modern languages define valid control flows, e.g., calls must always return to the instruction after the calling point, and many security-critical analyses such as data flow analysis rely on programs adhering to these valid control flows. However, today's compilation chains (compilers, linkers, loaders, runtime systems, hardware) cannot efficiently prevent violations of source-level control flows by co-linked low-level code, which can call and return to arbitrary instructions or smash the stack, blatantly violating the high-level abstraction. Other problems arise because languages fail to specify security policies, such as data confidentiality, and the compilation chains thus fail to enforce them, especially against powerful attacks such as those based on side channels. Yet other problems arise because enforcing source-level abstractions requires runtime checks with noticeable overhead, so compilation chains often forego security properties in favor of efficient code. The emerging field of secure compilation aims to address such problems by:

1. **Identifying precise security goals and attacker models.** Since there are many interesting security goals and many different kind of attacks to defend against, secure compilation is very diverse. Secure compilation chains may focus on providing (some degree of) type and memory safety for unsafe low-level languages like C and C++, or on providing mitigations that make exploiting security vulnerabilities more difficult. Other secure compilation chains use compartmentalization to limit the damage of an attack to only those components that encounter undefined behavior, or to enforce secure interoperability between code written in a safer language (like Java, C#, ML, Haskell, or Rust) and the malicious or compromised code it links against. Yet another kind of secure compilation tries to ensure that compilation and execution on a real machine does not introduce side-channel attacks.

2. **Designing secure languages.** Better designed programming languages and new language features can enable secure compilation in various ways. New languages can provide safer semantics, and updates to the semantics of old unsafe languages can turn some undefined behaviors into guaranteed errors. Components or modules in the source language can be used as units of

compartmentalization in the compilation chain. The source language can also make it easier to specify the intended security properties. For instance, explicitly annotating secret data that external observers or other components should not be able to obtain (maybe indirectly through side channels) may give the compilation chain the freedom to more efficiently handle any data that it can deduce is not influenced by secrets.

3. **Devising efficient enforcement and mitigation mechanisms.**
   An important reason for the insecurity of today's compilation chains is that enforcing security can incur prohibitive overhead or significant compatibility issues. To overcome these problems, the secure compilation community is investigating various efficient security enforcement mechanisms such as statically checking low-level code, compiler optimizations, software rewriting (e.g. software fault isolation), dynamic monitoring, and randomization. Another key enabler is the emergence of new hardware features that enable efficient security enforcement: access checks on pointer dereferencing (e.g. Intel MPX, Hardbound, WatchdogLite, Oracle SSM, SPARC ADI, or HWASAN), protected enclaves (e.g. Intel SGX, ARM TrustZone, Sanctum, or Sancus), capability machines (e.g. CHERI, Arm Morello), or micro-policy machines (e.g. Draper PUMP, Dover CoreGuard). The question is how such features can enable various security features in source languages efficiently, i.e., how hardware extensions can provide enforcement mechanisms for security properties.

4. **Developing effective verification techniques for secure compilation chains.**
   Criteria for secure compilation are generally harder to prove than compiler correctness. As an example, showing full abstraction, a common criterion for secure compilation, requires translating any low-level context attacking the compiled code to an equivalent high-level context that can attack the original source code. Another example is preservation of secret independent timing even in the presence of side-channels, as required for "constant-time" cryptographic implementations, which can require more complex simulation proofs than for compiler correctness. Finally, scaling such proofs up to even a simple compilation chain for a realistic language is a serious challenge that requires serious proof engineering in a proof assistant.

The Secure Compilation Dagstuhl Seminar 21481 attracted a large number of excellent researchers with diverse backgrounds. The 42 participants (12 on site, 30 remote) represented the programming languages, formal verification, compilers, security, systems, and hardware communities, which led to many interesting points of view and enriching discussions. Due to COVID-19 pandemic-related travel restrictions and uncertainties, many of the participants had to participate remotely using a combination of video conferencing, instant messaging, and ad-hoc gatherings. Despite this mixed environment, discussions thrived. Some of these conversations were ignited by the 5 plenary discussions and the 28 talks contributed by the participants. The contributed talks spanned a very broad range of topics: formalizing ISA security guarantees, hardware-software contracts, detection and mitigation of (micro-architectural) side-channel attacks, securing trusted execution environments, memory safety, hardware-assisted testing, sampled bug detection, formal verification techniques for low-level languages and secure compilation chains, machine-checked proofs, stack safety, integrating hardware-safety guarantees, effective compartmentalization and its enforcement, cross-language attacks, security challenges of software supply chains, capability machines, (over-)aggressive compiler optimizations, concurrency, new programming language abstractions, compositional correct/secure compilation, component safety, compositional verification, contextual and secure refinement, hardening WebAssembly, secure interoperability, (not) forking compilers, interrupts, hardware design, and many more. Talks were interspersed with lively discussions since, by default, each speaker could only use half of the time for presenting and had to use the other half for answering questions and engaging with the audience. Given the high interest spurred by this second edition and the positive feedback received afterwards, we believe that this Dagstuhl Seminar should be repeated in the future, when hopefully all the participants will be able to attend on site. One important aspect that could still be improved in future editions is spurring more participation from the systems and hardware communities, especially people working at the intersection of these areas and security or formal verification.

6.47 Representing and Solving Spatial Problems

Organizers: Pedro Cabalar, Zoe Falomir, Paulo E. Santos, and Thora Tenbrink

Seminar No. 21492

Date: December 5–10, 2021 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.11.11.1

Participants: Pedro Cabalar, Zoe Falomir, Werner Kuhn, Sabine Timpf, Florentin Wörgötter, Gabriel Zachmann

Remote Participants: Stefano Borgo, Anthony Cohn, Stefania Costantini, Simon Dobnik, Gavin Duffy, Anna Foerster, Ulrich Furbach, Shyamanta Moni Hazarika, Thiago Pedro Donadon Homem, David Kirsh, Felix Kroll, Pat Langley, Marcos Lopes, Bertram Ludäscher, Chayanika Deka Nath, Ana-Maria Olteteanu, Marianna Pagkratidou, Adam Pease, James Pustejovsky, Paulo E. Santos, Simon Scheider, Michael Sioutis, John Stell, Kristin Stock, Thora Tenbrink, Sabine Timpf, Maria Vasardani, Jago Williams, Diedrich Wolter

Preamble: The idea of organising a Dagstuhl Seminar on representing and solving spatial problems started with a discussion between Paulo Santos and the late Christian Freksa over a pair of Erdingers[46] during a previous Dagstuhl Seminar[47] in 2015. In this discussion we tried to solve, once and for all, the underlying reasoning strategies for spatial problem solving, whether it is logical-formal-symbolic, or it is totally/partially grounded on the real world (on what Christian called "the spatial substrates"). Considering the contributions in this volume, we believe the answer lies in the spectrum of approaches and that the right choice is dictated by the application domain. Indeed, one of the concrete contributions of this meeting was a discussion of possible problem scenarios that should constitute a Spatio-temporal Problem Repository to guide the future development of this field.

Summary of contributions: Spatio-temporal reasoning is a fundamental aspect of problem solving in general, and it occupies an ubiquitous place in AI and robotics, in the sense that almost all problem domains contain a spatio-temporal component. Paradoxically, despite its importance, the combination of spatial and temporal reasoning is one of the AI subfields with a wider gap between the current scientific progress and the human commonsense performance.

This seminar received contributions from researchers working in the various disciplines involved in investigating the problem of representing and reasoning about spatial problems (both in humans and artificial systems). The discussions were initially motivated by the following four main topics that are used below to organise the contributions in this collection:

1: **Representation**.
2: **Formalisation**.
3: **Description**.
4: **Problem solving**.

Representation

The discussions around the topic Representation were motivated by the following questions: How do humans conceptualise and mentally represent spatial problems? What is the role of high-level spatio-temporal structures for perceiving spatial

[46] One of which the former agent managed to spill all over the Schloss cellar floor!
[47] Dagstuhl Seminar 15192 – The Message in the Shadow: Noise or Knowledge? https://www.dagstuhl.de/en/program/calendar/semhp/?semnr=15192

problems, for manipulating spatial configurations, and for commonsense spatial problem solving?

In this volume, Cohn (Section 3.3) lists seven important problems that could be obstructing the development of automated commonsense spatial reasoning, including the lack of a proper definition of commonsense reasoning, the need for a foundational ontology of space, the problem of how to tackle vagueness and implicit knowledge and the need for suitable default reasoning mechanisms for dealing with spatial information. The acquisition of commonsense knowledge and the role of embodiment in perception and spatial awareness were also pointed out as key open research questions. Complementarily, Kuhn (Section 3.9) advocates for the following three conceptualisations for spatio-temporal phenomena: (i) space-time varying fields of attributes, (ii) objects and object collections, (iii) events over fields and objects. Wörgötter (Section 3.22) goes one step up in the abstraction level suggesting the investigation of the very idea of concepts and how stimulus-driven experience drives the formation abstract thought. Starting from a distinct standpoint, Sioutis' arguments (Section 3.17) go in a similar direction with the proposal of a generic neurosymbolic framework for integrating qualitative spation-temporal reasoning with neural models from a probabilistic perspective. The idea for a "(symbolic) spatio-temporal knowledge base, naturally grounded on physics and human cognition" could be a good starting point for testing representation frameworks. Langley (Section 3.10) describes a cognitive architecture for embodied agents, whose recent extensions (towards more flexible representation of space and processes) could provide a fertile ground for the development of research into the core problems in the representation and reasoning about commonsense spatial information.

Formalisation

This topic was motivated by the investigation of what would be a suitable formalism for commonsense problem solving that allows an accurate, flexible and readable knowledge representation for spatio-temporal effects of actions performed by an intelligent agent.

Much work in spatio-temporal reasoning has focused on *Qualitative Spatial Reasoning* (QSR) [11], a field that attempts the formalisation of spatial knowledge based on primitive relations defined over elementary spatial entities. Although the combination of (qualitative) temporal and spatial reasoning is not infrequent in the literature in general (see for instance [1]), QSR approaches have traditionally overlooked a formal treatment of actions. Aiming to bridge the gap between commonsense reasoning, reasoning about actions and change and qualitative spatial representation and reasoning, some recent research has focused on spatial puzzles and games [2, 3, 12], as these domains offer a small number of objects requiring a minimum background knowledge about unrelated features, while they keep enough complexity to constitute a challenging problem of KR. This is in line with Cabalar's contribution in this volume (Section 3.2), where the challenging problem of evidence analysis of digital forensics is proposed as a challenging domain for the application of spatio-temporal reasoning strategies. Calabar also suggests a set of minimum requirements for a KRR formalism about common sense, which includes simplicity, natural understanding, clear semantics, computability and elaboration tolerance, where logical formalisms allowing non-monotonic reasoning is pointed as the best candidate to fulfil these requirements.

One of the most modern computational tools for non-monotonic reasoning is Answer Set Programming (ASP). Ludäscher (Section 3.12) describes two technical challenges in the ASP encoding of QSR reasoning systems, namely the difficulty in distinguishing spatial configurations using the natural ASP encoding of pairwise relations; and the exploration of the hierarchical structure of various QSR formalisms. While Ludäscher describes the latter issue from an implementation perspective, Stell (Section 3.18) suggests a research program to investigate the interaction with hierarchical models of discrete space and time representing changes at different levels of detail. The key idea here is to combine granular change with temporal change for supporting reasoning about effects of actions at different levels of detail.

From a more foundational standpoint, Borgo (Section 3.1) argues for the creation of a systematic program for the formal study of spatial notions, such as between, convex, simplex, parallel etc. This program should include the comparison of distinct perspectives for such modelling, facilitating the understanding of cognitive representation and reasoning. The notion of path is given as an example of how this systematic program could be developed.

Pease (Section 3.14) proposes the formalisation of everyday concepts and facts in a high-level classical logic engine called Suggested Upper Merged Ontology (SUMO). This formalisation would constitute a collection of real-world spatio-temporal problems that could be used to test the sufficiency of spatial knowledge representation strategies.

Description

This topic deals with the development of human readable descriptions of the inputs, reasoning steps and solutions of spatial problems. In particular, it addresses whether (and to what extent) it would be possible to develop high-level representations or interfaces for dealing with natural language and/or diagrammatic descriptions that allow specifying both the input knowledge and the output conclusions in terms of textual descriptions of spatial problems.

Research on spatio-temporal language highlights the range of meanings across contexts [13] as well as patterns of usage in relation to mental representations [14] and problem-solving [15] in various domains. Part of the insight gained in this research concerns the significance of what can (or will normally) *not* be represented in language. Often, non-verbalised concepts are those that are best understood through features of the spatial world itself. For instance, we may verbalise only the high-level goal of an everyday action (such as *dress up*), because every detailed action is represented through a combination of world knowledge with the affordances of the actual objects in question. There is no need to learn or conceptualise how to handle every possible instance of these objects, because the affordances of each exemplar are sufficiently clear to act upon as required. Some of these issues feature in the contributions of Dobnik (Section 3.4), Lopes (Section 3.11), Scheider (Section 3.16), Stock (Section 3.19) and Tenbrink (Section 3.20)

The following three semantic representational dimensions are listed by Dobnik (Section 3.4), as a summary of the core aspects of spatial descriptions studied in areas of linguistics, psychology, and computer science: (i) scene geometry, (ii) world knowledge and (iii) dynamic aspects of interaction. Dobnik's contribution also points out to a few open problems in this field, such as the lack of a unified computational language model that includes all of the cited dimensions, the limitation in the current research of grounding language in perception, that mostly deals with geometric perceptual context, and the need for experimental evaluation of models in real open domains.

Lopes (Section 3.11) suggests the use of the differences and interrelations between maps and navigation instructions to instan-

tiate three cognitive core strategies of problem solving (attributed independently to McCarthy and Sloman), which are solving problems by following instructions, by descriptive knowledge representation and by analogical reasoning. Complementary, Stock (Section 3.19) argues that the current modelling of relative location descriptions is too restricted, not taking into account important contextual factors such as physical environment and objects in it; the observer's goals and expectations; the audience location and knowledge. Two main research advances were then proposed to push the boundaries of this field: (i) the incorporation of factors proved to be important in linguistic and cognitive science research and (ii) cognitive science investigation on some neglected factors that could be used in computational models, such as the influence of users' goals in the use of spatial relation terms. Tenbrink (Section 3.20) discusses the gap between human descriptions of routes and the state-of-the-art in the automatic generation of such descriptions to motivate three main challenges that deserve further attention: the combination of visual and verbal information, the consideration of change over time, and the flexibility in the use of various reference systems. Scheider (Section 3.16) contributes to this discussion suggesting the study of *transformation models* for concepts that are constitutive of spatial information, such as the concept of location, field, object, event and network. Arguing that transformations lie at the core of spatial reference systems, the author suggests that transformation models could be used to handle the variability of spatial information conceptualisations.

Problem Solving

The discussion about problem solving was initially motivated by questions about whether or not it would be possible (and desirable) to develop interfaces for dealing with spatial configurations including diagrammatic depictions and natural language descriptions to solve spatial puzzles in similar ways as humans do; and also, what are the commonsense problem solving capabilities involving spatio-temporal features including temporal explanation and planning under physical/geometric qualitative or semi-quantitative constraints. This issue also includes the investigation of appropriate problem solving algorithms and their potential applications to real-world domains that could be of interest to industry.

Calculi for qualitative temporal and spatial reasoning have a long history [4,5]. More recently the relation between 2D and 3D structures of space (plus time) and their 1D formal descriptions have been investigated; the goal is a farther reaching exploitation of spatial structure for spatial problem solving. Cognitive architectures for robotic agents that make direct use of knowledge and spatial affordances in physical environments ('knowledge in the world') are currently being developed [10]. The knowledge in the mind of these agents controls their perception and action in ways that simplifies spatial configurations in order to reduce the difficulty of solving the spatial problem at hand, while knowledge about spatial relations is represented directly in perceivable and manipulable physical structures [9]. A principle underlying this approach is the use of mild abstraction [7] as employed in map navigation or in constructive geometry [8]. Much QSR and related reasoning research was originally motivated by the obvious discrepancy between logic-based metric computation (where spatial directions may be defined, for instance, by exact geometric angles and location specifications) and human concepts (such as "to the left"). Falomir's contribution (Section 3.5) starts from a discussion of the need to use cognitive heuristics for problem solving directly acting in the physical world, and argues for an effective combination of machine learning with automated reasoning for spatial problem solving. Spatial reasoning problems based on Perceptual/Differential Ability tests (PAT/DAT) are suggested as test domain for this development. Instead of assuming PAT/DAT as test domain, but still advocating the investigation of cognitive heuristics, Kroll (Section 3.8) describes recent results on the use of gaze heuristics for guiding an autonomous agent, and suggests the use of these ideas to dexterous assembly tasks; whereas Zachmann (Section 3.23) describes several issues pertaining the development of 3D geometric simulation methods for discovering affordances.

Hazarika (Section 3.6) discusses the use of mental maps defined over the space-time information structures underlying spatial problems (the *Spatial Substrate* [6]) for understanding how objects and their affordances affect the way humans reason about space. Following a similar line, Nath (Section 3.13) suggests the use of diagrammatic representations and reasoning to express the spatial problems and their physical substrate.

More explicit representation tools for problem solving were described by Homem (Section 3.7) and Santos (Section 3.15), in which they propose extending their previous work on applying case-based reasoning combined with QSR techniques towards complex real-world problems, such as real-time strategy games (Section 3.7) and collaborative mission planning for autonomous maritime vehicles (Section 3.15). Similarly, Wolter (Section 3.21) advocates the investigation of spatial problem solving techniques in dynamic domains (e.g. physical manipulation games) that allow for the progressive development of systems towards human-level capabilities.

Spatio-temporal Problem Repository

An issue that was of common agreement between the participants of the seminar is the need to create a Spatial Reasoning Problem repository. A few possible domain scenarios were discussed during the seminar, as listed below, but their full description is left for a future document.

1. Digital Forensics, cited in Section 3.2;
2. Puzzles such as Crazy Machines and Cut the Rope, cited in Section 3.2;
3. Angry birds, cited in Section 3.21
4. TPTP-style spatial problems, cited in Section 3.14
5. Perceptual/Differential Ability tests (PAT/DAT), cited in Section 3.5
6. Interpretation and generation of map descriptions (discussed during the seminar)
7. Multi step reasoning VQA (discussed during the seminar)
   - Ignore the images and concentrate on the scene graphs
   - Turn the scene graphs into logical formalism
8. From puzzle's descriptions to solutions (discussed during the seminar)
   - Input: Puzzles described as diagrams, or described with words or both
   - Output: A solution
   - it involves: Multiple step reasoning, combining diagrams, language and problem solving

References

1 Bennett, B., Cohn, A. G., Wolter, F., and Zakharyaschev, M. (2002). Multi-dimensional modal logic as a framework for spatio-temporal reasoning. *Applied Int.*, 17:239–251.

2 Cabalar, P. and Santos, P. E. (2011). Formalising the fisherman's folly puzzle. *Artificial Intelligence*, 175(1):346–377.

3 Cabalar, P. and Santos, P. E. (2020). Spatial reasoning about string loops and holes in temporal ASP. In Calvanese, D., Erdem, E., and Thielscher, M., editors, *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning, KR 2020, Rhodes, Greece, September 12-18, 2020*, pages 182–192.

4 Freksa, C. (1991). *Cognitive and linguistic aspects of geographic space*, chapter Qualitative spatial reasoning. Kluwer, Dordrecht.

5 Freksa, C. (1992). Temporal reasoning based on semi-intervals. *Artificial Intelligence*, 54(1):199 – 227.

6 Freksa, C. (2015). Computational problem solving in spatial substrates – A cognitive systems engineering approach. *Int. J. Softw. Informatics*, 9(2):279–288.

7 Freksa, C., Barkowsky, T., Dylla, F., Falomir, Z., Olteteanu, A.-M., and van de Ven, J. (2018a). Spatial problem solving and cognition. In Zacks J, T. H., editor, *Representations in Mind and World*, pages 156–183. New York: Routledge.

8 Freksa, C., Barkowsky, T., Falomir, Z., and van de Ven, J. (2018b). Geometric problem solving with strings and pins. *Spatial Cognition & Computation*. (in press).

9 Freksa, C., Olteteanu, A.-M., Ali, A. L., Barkowsky, T., van de Ven, J., Dylla, F., and Falomir, Z. (2016). Towards spatial reasoning with strings and pins. In *Fourth Annual Conference on Advances in Cognitive Systems*, volume 4.

10 Freksa, C., Vasardani, M., and Kroll, F. (2020). Dynamic problem solving in space. In Šķilters, J., Newcombe, N. S., and Uttal, D., editors, *Spatial Cognition XII*, pages 18–32, Cham. Springer International Publishing.

11 Ligozat, G. (2013). *Qualitative Spatial and Temporal Reasoning*. John Wiley & Sons.

12 Santos, P. E. and Cabalar, P. (2016). Framing holes within a loop hierarchy. *Spatial Cognition & Computation*, 16(1):54–95.

13 Tenbrink, T. (2007). *Space, time, and the use of language: An investigation of relationships*. Mouton de Gruyter, Berlin.

14 Tenbrink, T., Andonova, E., Schole, G., and Coventry, K. (2017). Communicative success in spatial dialogue: The impact of functional features and dialogic strategies. *Language and Speech*, 60:318–329.

15 Tenbrink, T. and Taylor, H. A. (2015). Conceptual transformation and cognitive processes in origami paper folding. *Journal of Problem Solving*, 8:2–22.

7 Öffentlichkeitsarbeit
Public Relations and Outreach

7.1 Pressemitteilungen und Medienarbeit / Press Releases and Media Work

Die regelmäßige Erstellung und Herausgabe von Pressemitteilungen dient der verständlichen Verbreitung von aktuellen Informatikthemen. Die Vermittlung des Konzepts von Schloss Dagstuhl ist dabei ebenfalls ein Thema. Pressemitteilungen und Berichterstattungen in diversen Medien – soweit bekannt – sind über das Internetportal von Schloss Dagstuhl[48] 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. 2 660 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. 1 850 Abonnenten verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ gepflegt (derzeit etwa 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[48].

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 2,660 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 1,850 followers. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with about 620 members), which supports the networking of participants in Dagstuhl Seminars. Additionally, it also serves as a forum for announcing interesting news about Schloss Dagstuhl.

[48] https://www.dagstuhl.de/about-dagstuhl/press/

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 institute 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 (außerhalb von Pandemiezeiten) 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 in summer and two weeks in 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 (in non-pandemic times) 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 auch für Hybridveranstaltungen ausgestattet. Darunter verstehen wir Veranstaltungen mit Teilnehmern die vor Ort teilnehmen und anderen, die per Videokonferenztechnik von ihrem perönlichen Rechner aus oder mittels der Technik eines der anderen Hörsäle teilnehmen. Die Technik dafür umfasst eine Audioanlage mit Lautsprechern und verschiedenen Mikrophonen. Neben einem Kopfbügelmikrophon und Handmikrophon sind mehrere Mikrophone an den Decken montiert, sodass alle Teilnehmer frei diskutieren können. Für die Videoübertragung stehen eine fernsteuerbare Kamera in Richtung des Vortragenden und der Tafeln sowie eine fest ausgerichtete Kamera auf das Publikum zur Verfügung. Weiterhin sind jeweils ein zentraler Beamer für die Präsentation und ein weiterer Beamer bzw. ein 85 Zoll Display zur Ansicht der Remote-Teilnehmer vorhanden. Daneben steht in jedem Hörsaal auch ein MS-Windows-Laptop zur Verfügung.

Conference Facilities

Schloss Dagstuhl has three lecture halls with a seating capacity of 25 to 60 people each. All lecture halls are equipped for hybrid events. This entails that offered events can include a combination of on-site participants and those participating via video conference from their own personal computer or from one of the other lecture halls. The technology offered includes an audio system with speakers and various microphones. In addition to a headset microphone and a hand-held microphone, several microphones are mounted on the ceiling to foster discussions between all participants. For the video transmission, we offer a remote-controlled camera for the speaker and the content of the boards, as well as a fixed camera capturing the on-site audience. Furthermore, there is a projector for the speakers presentation and an additional projector or an 85 inch display showing the remote participants. A laptop with MS-Windows is available in every lecture hall.

In addition to the lecture halls, the center has six meet-

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, selbst der Weinkeller bietet ein Whiteboard. Weitere mobile Flipcharts stehen flexibel für alle Veranstaltungen bereit.

Daneben gibt es über das ganze Zentrum verteilt weitere Räume, in denen Gäste sich in entspannter Atmosphäre treffen und diskutieren können. Insbesondere am Abend zieht es viele Gäste in den Weinkeller und die Cafeteria, zwei der gemütlichsten Räume im Haus und hervorragend geeignet für die Fortsetzung einer produktiven Diskussion in angenehmer Atmosphäre.

Dagstuhls Küche

Die Mahlzeiten sind ein wichtiger Bestandteil des wissenschaftlichen Programms von Schloss Dagstuhl. Die Sitzordnung wird absichtlich stets zufällig gemischt, um eingefahrene Gruppen aufzuteilen und Gäste zu ermuntern, während ihres Aufenthalts möglichst viele verschiedene Kollegen kennenzulernen. Große Tische im Speiseraum fördern die gemeinschaftliche Interaktion bei den Mahlzeiten.

Dagstuhls Philosophie des Kochens ist einfach: saisonal, gesund und schmackhaft. Unsere Gerichte werden jeden Tag von unseren 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, mitgeing 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. Even the wine cellar offers a whiteboard. Additional mobile flipcharts are available for all events.

The center also offers a variety of other spaces where guests can sit and work together in a relaxed atmosphere. Particularly in the evening, guests gravitate towards the wine cellar and upstairs café, two of the coziest places in the house and great places for continuing a productive discussion in a comfortable atmosphere.

Dagstuhl's Kitchen

The dining experience at Dagstuhl is an important part of the center's scientific program. Seating arrangements are mixed deliberately in order to break up cliques and encourage guests to talk to as many different people as possible during the course of their stay. Large tables in the dining hall promote collaborative interaction during meals.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy, and tasty meals. Everything is freshly prepared each day by the kitchen's staff. 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 prior 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 people, opens onto the castle garden and patio, and offers a relaxed,

brachte 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.

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 müssen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser von Schloss Dagstuhl subventionierte 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 prinzipiell gut genutzt. Pandemiebedingt und insbesondere auch durch die Möglichkeit der Remoteteilnahme für Teilnehmer der meisten Seminare, beherbergte Schloss Dagstuhl 2021 nur ein Kind, das von den Eltern betreut wurde.

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

familiar atmosphere.

Small and late-morning breaks break up 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.

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 is subsidized by Schloss Dagstuhl, can be booked for a small fee prior to the seminar.

Parents also have the option to bring along a caregiver of their own, usually a spouse or relative. In the case of seminar participants, the children’s costs for room and board are covered by Schloss Dagstuhl. If instead of using Dagstuhl’s childcare service, a caregiver accompanies the child or children, the cost for room and board for the caregiver is also covered by Schloss Dagstuhl.

Dagstuhl’s childcare offer for parents is in general a well used service. Due to the pandemic and in particular due to the option of remote participation for participants of most seminars, Schloss Dagstuhl hosted only one child in 2021, who was cared for by the child’s parents.

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

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.

instruments, e.g., two concert guitars. Schloss Dagstuhl also has a full sauna, a pool table, table football facilities, mountain bikes, a dartboard, and a recreation room with gym equipment and table tennis as well as outdoor sports grounds featuring a volleyball net.

Geschäftsstelle in Saarbrücken 8.2

Dagstuhl Office 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.

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.

Dagstuhl an der Universität Trier 8.3

Dagstuhl at the University of Trier

Die für die Bibliographiedatenbank dblp zuständigen Mitarbeiter haben ihren Standort an der Universität Trier. Die Ende 2010 zunächst auf Basis zweier Projekte gestartete Zusammenarbeit zwischen Schloss Dagstuhl und der Universität Trier wurde im November 2018 in eine offizielle und permanente Außenstelle von Schloss Dagstuhl auf dem Campus der Universität Trier überführt. Dabei profitiert das dblp-Team von der engen Zusammenarbeit mit der Abteilung Informatikwissenschaften und als externer Partner im Digital Research and Bibliographic Meta Data Lab des Center for Informatics Research and Technology (CIRT).

The scientific and editorial staff working on the *dblp computer science bibliography* is located at the Dagstuhl offices at the University of Trier. Initially based on a project-based cooperation between Schloss Dagstuhl and the University of Trier which was first established in 2010, in November 2018, an official and permanent Schloss Dagstuhl branch office has been established on the campus of the University of Trier. In Trier, the dblp team benefits from the close cooperation with the University's department of computer science, 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.

Bibliothek

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 36 000 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 Springer Nature spendet der Bibliothek alle Bände der Reihe Lecture Notes in Computer Science (LNCS) in Druckform. Die Bibliothek verfügt somit über Druckexemplare aller veröffentlichten Bände ab dem Band 1.
- Wissenschaftliche Fachzeitschriften sind eine wesentliche Voraussetzung für exzellente Forschung. Häufig werden in Zeitschriften erweiterte Fassungen von

Schloss Dagstuhl has two central services: the IT service and a research library, which are both located at the main site in Wadern.

9.1 *Research Library*

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 36,000 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 printed volumes of the series Lecture Notes in Computer Science (LNCS) to the library. 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 licenses, as well as consortium licenses funded by the

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.

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 komplette Monographienbestand 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.

Schloss Dagstuhls Fachbibliothek ist institutionelles Mitglied des Deutschen Bibliotheksverbandes (DBV). Die Bibliothekarin Petra 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 complete Monograph collection 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.

Schloss Dagstuhl's specialized library is an institutional member of the German Library Association (Deutscher Bibliotheksverband, DBV). The librarian, Petra 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 zahlreichen Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2021 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 598 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. In 2021, 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 598 volumes during the year 2021 as donations from publishing houses and seminar participants.

Fig. 9.1
Donations from publishers to the Dagstuhl library.

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[49] (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 Schloss Dagstuhls Publikationsplattform DROPS (https://drops.dagstuhl.de) befindet und
- den dblp-Server (https://dblp.dagstuhl.de oder, alternativ, 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[49] (which is a comfortable option for guests with existing *eduroam* accounts). Within its facilities, Schloss Dagstuhl provides a vast 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 photocopier.
- 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 https://drops.dagstuhl.de, Schloss Dagstuhl's publishing platform, and
- the dblp server at https://dblp.dagstuhl.de or, alternatively, 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.

[49] *eduroam* (education roaming) is a world-wide roaming access service developed for the international research and education community, see https://www.eduroam.org.
eduroam *(education roaming) ist ein weltweiter roaming Zugangsdienst der für die internationale Forschungs- und Bildungsgemeinde entwickelt wurde, siehe https://www.eduroam.org.*

10 Kunst
Art

Dagstuhl als Galerie

10.1

Dagstuhl as an 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. Im Ergebnis wurden bis 2018 insgesamt drei Ausstellungen aus dem Fundus von Saartoto zusammengestellt und in Dagstuhl präsentiert.

Pandemiebedingt fand in 2021 leider keine Ausstellung statt. 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 recently, the exhibitions were organized by artists and individual collectors. The year 2016, however, saw the establishment of a cooperation between Saarland-Sporttoto GmbH (Saartoto for short), Hochschule für Bildende Künste Saar (university of art and design; HBKsaar for short), and Schloss Dagstuhl, which makes Saartoto's collection accessible to Schloss Dagstuhl for a series of exhibitions. Being a major art sponsor, Saartoto is in possession of a substantial art collection. In the context of this collaboration, HBKsaar takes stock of and documents Saartoto's art collection. As a result, a total of three exhibitions from the Saartoto collection had been curated and presented in Dagstuhl until 2018.

Due to the Covid-19 pandemic, unfortunately, there was no exhibition hosted by Schloss Dagstuhl in 2021. Current exhibitions are open to the interested public upon request.

Dagstuhls permanente Kunstausstellung

10.2

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 ausgewählten Ausstellungen 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, an artwork from selected exhibitions is chosen and donors are drummed up. Thus, approximately 180 works of art have been acquired over the last 25 years. Additionally, this initiative has increasingly encouraged artists and patrons to make donations. The artworks adorn the rooms of Schloss Dagstuhl in Wadern as well as the Dagstuhl Office in Saarbrücken.

11 Struktur der Gesellschaft
Structure of the Company

Gründung und Gesellschafter 11.1 *Formation and Shareholders*

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit elf Gesellschaftern (siehe Fig. 11.1) organisiert. Dies sind 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. (GI),[50] 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[51] (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”[52] to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”[53] 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 was 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.

[50] engl.: German Informatics Society
[51] engl.: Federal/State Government Commission for Educational Planning and Research Promotion
[52] engl.: International Conference and Research Center for Computer Science
[53] 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 ist Prof. Raimund Seidel, Ph. D., der wissenschaftliche Direktor von Schloss Dagstuhl.

Shareholders' Meeting

All changes to the company, in particular the inclusion of new associates, the revision of the Shareholders' agreement, and the dissolution of the company, are decided by the shareholders. Shareholders also confirm new members forwarded by them to the Supervisory Board and the appointment or recall of the managing directors. In accordance with their shares, all shareholders currently have the same number of votes except the Gesellschaft für Informatik, which has three times the number of votes of 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.

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee

Der Wissenschaftliche Direktor wird dem Aufsichtsrat von einer Findungskommission zur Berufung vorgeschlagen. Dieser Findungskommission gehören mindestens der Vorsitzende des Aufsichtsrats und der Vorsitzende des Wissenschaftlichen Beirats an. Die Amtszeit des Wissenschaftlichen Direktors beträgt fünf Jahre.

Die technischen und administrativen Aufgaben werden vom Technisch-administrativen Geschäftsführer wahrgenommen. Seit Juli 2014 hat Heike Meißner diese Position inne.

consisting of at least the chairperson of the Supervisory Board and the chairperson of the Scientific Advisory Board. The term of office of the Scientific Director is five years.

The Technical Administrative Director is responsible for technical and administrative tasks. Since July 2014, 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 a 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 is three years – with the exception of the Scientific Director. 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 Wissenschaftliche Beirat auf zwei Wegen bei der Qualitätssicherung mit. Zum einen berät er die Leitung in Fragen der

Scientific Advisory Board

The tasks of the Scientific Advisory Board (see Fig. 11.5) are not only defined by the Shareholders' Agreement, but also by the recommendations of the Leibniz Association. The latter stipulates two different ways in which the Scientific Advisory Board is involved in quality assurance. On the one hand, the board offers advice to the

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.

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 the making of important decisions with regard to the 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. | *tenure ended in December 2021*
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany | Representative of Karlsruher Institut für Technologie | *tenure started in October 2021*
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
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, Germany | 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 | *tenure ended in October 2021*
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
Prof. Dr.-Ing. Franz Baader TU Dresden, Germany
Prof. Gilles Barthe, Ph. D. IMDEA Software Institue, Madrid, Spain and Max Planck Institute for Security and Privacy, Bochum, Germany | *tenure ended in October 2021*
Prof. Dr. Daniel Cremers Technische Universität München, Germany
Goetz Graefe Google Inc., Madison, United States of America
Prof. Dr. Reiner Hähnle TU Darmstadt, Germany
Prof. Dr. Barbara Hammer Universität Bielefeld, Germany
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
Dr. Steve Kremer Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy – Grand Est, France
Prof. Rupak Majumdar, Ph. D. Max Planck Institute for Software Systems, Kaiserslautern, Germany | *tenure started in November 2021*
Prof. Dr.-Ing. Heiko Mantel TU Darmstadt, Germany | *tenure started in November 2021*
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany | *tenure ended in October 2021*
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
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 | Chair of the Scientific Advisory Board
Prof. Dr. Anja Feldmann Max-Planck-Institut für Informatik, Saarbrücken, Germany
Prof. Dr. Ir. Dr. h. c. (AAU) Joost-Pieter Katoen RWTH Aachen, Germany
Prof. Dr. Friedhelm Meyer auf der Heide Heinz Nixdorf Institute, Paderborn and Universität Paderborn, Germany | *tenure ended in May 2021*
Prof. Dr. Laurent Romary Centre Marc Bloch, Berlin, Germany and INRIA, Paris, France
Prof. Alistair Sinclair, Ph. D. University of California, Berkeley, United States of America | *tenure will start in January 2022*
Prof. Dr. Lothar Thiele ETH Zürich, Switzerland | *tenure will start in January 2022*

Fig. 11.5
Scientific Advisory Board.

Kuratorium | Industrial Curatory Board
Dr. Tim Harris Microsoft Research, Cambridge, United Kingdom
Dr. Jaroslaw Kutylowski DeepL, Köln, Germany | *tenure will start in January 2022*
Christine Regitz SAP SE, Walldorf, Germany
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“

Association “Friends of 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 ist mit Hilfe einer professionellen und auf Stiftungen spezialisierten Vermögensverwaltungsgesellschaft angelegt und entwickelt sich auch in Zeiten hoher Dynamik insgesamt gut.

Der Verein hat im vergangenen Jahr Schloss Dagstuhl inbesondere durch den Ankauf von Kunstwerken unterstützt, und fördert den längerfristigen Plan, die Durchführung von Ausstellungen auf eine institutionelle Basis zu stellen.

In Folge der insgesamt positiven Erfahrungen, die ursächlich der Pandemie geschuldet waren, hat der Verein zuletzt eine Satzungsänderung beschlossen, die es erlaubt auch weiterhin Mitgliederversammlungen des Vereins per Videokonferenz durchzuführen. Dies dient dazu, einen intensiven Austausch aller Mitglieder zu ermöglichen, ohne dabei lange Anreisewege und -zeiten zu bedingen.

Weitere Informationen zum Verein, aber auch Mitgliedschaftsanträge finden Sie unter http://www.friends-of-dagstuhl.de.

Since mid 2014, the registered association for 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, the 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 have been invested with the help of a professional asset management company specialising in foundations. The assets develop well overall even in times of high volatility.

Last year, the association supported Schlosss Dagstuhl in particular through the purchase of works of art, and is supporting the longer-term plan to put the organisation of exhibitions on an institutional basis.

As a result of the overall positive experience gained during the pandemic, the association recently decided to amend its statutes to allow it to continue to hold general meetings of the association via video conference. The purpose of this is to enable an intensive exchange between all members without requiring long journeys and travel times.

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) Saarbrücken, Germany
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. Durch die Covid-19 Pandemie bedingt, weichen die veranstaltungs- und teilnehmerbezogenen statistischen Werte der Jahre 2020 und 2021 erheblich von denen der vorhergehenden Jahre bis einschließlich 2019 ab. Details sind in den Abschnitten weiter unten 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. Die sehr hohe Anzahl von Anträgen 2020 kann unter anderem auf die pandemiebedingte auf Juni 2020 verschobene zweite Antragsfrist zurückgeführt werden. In dieser Runde wurde mit 92 Anträge soviele wie nie zuvor gestellt. Die sehr niedrige Anzahl 2021 ist ausschließlich durch eine beschränkte erste Antragsrunde, in der nur wegen der Pandemie abgesagte Seminare neu beantragt werden durften, begründet. In dieser Runde wurden schließlich alle Anträge, die ja schon einmal positiv begutachtet worden waren, genehmigt, was den überdurchschnittlich hohen Anteil der genehmigten Seminare erklärt. Fig. 13.2 zeigt, wie die akzeptierten Seminare und Workshops sich bezüglich Größe und Länge aufgliedern.

Veranstaltungs-bezogene Daten: Fig. 13.3 zeigt die Anzahl der verschiedenen Veranstaltungstypen. Der dramatische Einbruch 2020 und die immer noch geringe Anzahl von Veranstaltungen in 2021 sind ebenso durch die Covid-19 Pandemie bedingt. In 2020 hatte Schloss Dagstuhl wegen der Pandemie ab Mitte März für etwa 5 Monate geschlossen. Aber auch nach Wiedereröffnung sind 2020 fast alle geplanten Veranstaltungen abgesagt worden. Beginnend mit 2021 bietet Schloss Dagstuhl übergangsweise auch hybride Seminare, in denen Teilnehmer sowohl vor Ort sind als auch online über Audio-/Videoübertragung zugeschaltet sind, an. Als Spezialfall sind in den reinen online Seminare gar keine Teilnehmer vor Ort. Alle diese Typen werden in der Tabelle hier mitgezählt. Fig. 13.4 gibt aber einen Überblick, wie viele Veranstaltungen on-site, hybrid oder online stattgefunden haben.

Daten zu der Anzahl der durchgeführten Seminare gegliedert nach der ursprünglich genehmigten Größe und Dauer sind in Fig. 13.5 angegeben. Manche Organisatorenteams von – insbesondere rein online stattfindenden – Seminaren hatten sich entschlossen, die Dauer teilweise erheblich zu verkürzen.

Zum Ausgleich der Einschränkungen von online oder hybrid stattfindenden Seminaren wurden bei diesen auch mehr Teilnehmer als ursprünglich genehmigt zugelassen. Fig. 13.6 zeigt Anzahl und Anteil der eingeladenen Seminarteilnehmer, welche die Einladung

This section provides statistical data about the scientific program and the composition of program participants. Due to the Covid-19 pandemic, the statistical values related to the program and participants of 2020 and 2021 deviate considerably from those of previous years up to and including 2019. Details are provided in the sections below.

Charts and tables are structured as described below.

Proposal-related data: Fig. 13.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops as well as the respective acceptance rates. The exceptionally large number of proposals in 2020 is due to the fact that the second submission deadline was postponed until June 2020, among other reasons. That proposal round saw 92 submissions, more than ever before. The very low number in 2021 is exclusively due to the fact that the first submission round was limited in that only resubmissions of proposals for seminars that had been cancelled due to the pandemic were allowed. Eventually, all proposals were accepted – they had obviously been accepted before – which explains the above-average share of accepted seminars. Size and duration of accepted seminars and workshops are displayed in Fig. 13.2.

Event-related data: Fig. 13.3 illustrates the number of different event types. The sharp drop in 2020 and the continuously low number of events in 2021 were also caused by the Covid-19 pandemic. From mid-March 2020, Schloss Dagstuhl was closed for approximately 5 months as a consequence of the pandemic. Even after re-opening, almost all scheduled events in 2020 were cancelled. Beginning in 2021 and on an interim basis, Schloss Dagstuhl has been offering hybrid seminars where some participants are on site while others join the seminar online via audio/video transmission. In special cases, there are also online-only seminars where there are no participants on site. The table includes all of these event types. However, Fig. 13.4 provides an overview of how many events took place on site, in a hybrid setting or online, respectively.

Fig. 13.5 illustrates data regarding the number of seminars that took place, categorised by originally approved size and duration. Some teams of seminar organizers – especially in the context of online-only seminars – decided to shorten the duration, in some cases drastically.

In order to compensate for the limitations of hybrid and online seminars, larger numbers of participants than originally approved were allowed to join. Fig. 13.6 shows the number and proportion of invited seminar participants who accepted or declined the invitation. It is evident that the share of invitees who accepted their invitation was significantly higher in 2021. This may

annehmen bzw. ablehnen. Der in 2021 auffallende, deutlich höhere Anteil an Eingeladenen, die zugesagt haben, mag zum einen an der Möglichkeit der Online-Teilnahme eines Teils der angebotenen Seminare liegen, aber genauso auch daran, dass zu manchen Seminaren noch sehr kurzfristig vor dem Seminar Personen eingeladen wurden, deren Zusage vor der Einladung abgesprochen wurde. Auch dies wurde, zum Teil, durch die Möglichkeit der Online-Teilnahme ermöglicht.

Die Verteilung der Annahmerate pro Seminar, also der Quotient aus Teilnehmer und allen Eingeladenen, ist in Fig. 13.7 dargestellt. Durch die pandemiebedingten Besonderheiten des Seminarbetriebs in und den deutlich kleineren Stichproben aus 2020 und 2021 sind die Daten dieser Jahre nicht im Vergleich der anderen Jahre zu interpretieren.

Fig. 13.8 zeigt dagegen, wie viel Prozent der zugesagten Größe (gemessen an der Personenanzahl) tatsächlich von einem Seminar belegt wurde. Auch hier sind die Daten der letzten beiden Jahre nicht mit den vorherigen Daten zu vergleichen. Werte über 100 % bedeuten, dass mehr Personen als ursprünglich genehmigt am Seminar teilgenommen haben. Während dies bei online zugeschalteten Teilnehmer problemlos möglich ist, war die Gesamtanzahl der Teilnehmer im Haus und damit auch für die Veranstaltungen strikter limitiert.

Teilnehmer-bezogene Daten: Grundsätzlich werden hier alle Teilnehmer gezählt, unabhängig davon, ob sie vor Ort oder online teilgenommen haben. Vor allem durch ausfallende Veranstaltungen aber auch durch schlechtere Beteiligung insbesondere an Veranstaltungen die ausschließlich vor Ort stattfanden, sind die Zahlen aus 2020 und 2021 deutlich geringer als die der Vorjahre.

Die Teilnehmerzahlen – abhängig vom Veranstaltungstyp – gibt Fig. 13.9 an. Einen Eindruck, wieviele Teilnehmer vor Ort waren bzw. remote teilgenommen haben, vermittelt Fig. 13.10.

Fig. 13.11 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. 13.12 entnommen werden. Die Anzahl von früheren Seminarbesuchen kann man Fig. 13.13 entnehmen. Fig. 13.14 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. 13.15 gezeigt.

Auslastungs-bezogene Daten: Die Auslastung des Zentrums wird schließlich in Fig. 13.16 an Hand der Übernachtungen und ihrer Verteilung über die einzelnen Wochen getrennt nach Veranstaltungstypen aufgezeigt. Im Gegensatz zu den anderen Statistiken beziehen sich diese Daten ausschließlich auf Gäste vor Ort.

in part be due to the fact that some seminars offered online participation, but may also have been caused by issuing invitations to some seminars at very short notice when the invitees had informally agreed to attend in advance. This was also partly made possible through offering online participation.

The acceptance rate distribution per seminar, i.e., the quotient of the number of participants and the total number of invitees, is illustrated in Fig. 13.7. Due to the particularities regarding the seminars and the significantly smaller samples in 2020 and 2021 caused by the pandemic, the data from those years is not to be interpreted in relation to previous years.

In contrast, Fig. 13.8 visualizes the percentage of the reserved space (in terms of number of people) that was actually used by seminar participants. Again, the data from the last two years cannot be compared to previous data. Values above 100 percent show that more people than originally approved attended the seminar. While this was easy to facilitate with regard to online participants, the total number of participants on site and therefore the maximum number of on-site participants allowed for the respective events was limited more strictly.

Participant-related data: Basically, all participants are counted here, regardless of whether they participated on site or online. The numbers from 2020 and 2021 are significantly lower than the ones in previous years due especially to event cancellations, but also due to lower participation rates, particularly with regard to on-site only events.

Fig. 13.9 shows the number of participants broken down by event type. Fig. 13.10 illustrates how many participants were on site and how many participated remotely.

Fig. 13.11 shows the distribution of our guests' country affiliations.

Survey-related data: In this section we present selected data obtained from our ongoing Dagstuhl Seminar and Dagstuhl Perspectives Workshop guest survey project. An overview of the results gained from routine participants surveys for Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Fig. 13.12. Fig. 13.13 displays how often participants have attended seminars in the past. Fig. 13.14 provides data on the seniority of participants. While Dagstuhl Seminars and Dagstuhl Perspectives Workshops mainly target academic researchers, researchers and developers from industry are also always welcome. The distribution of their share is shown in Fig. 13.15.

Utilization-related data: Finally, Fig. 13.16 illustrates utilization on the basis of overnight stays hosted at Schloss Dagstuhl – broken down by event type – as well as their distribution by week. In contrast to the other statistics, this data exclusively refers to guests on site.

Gender-related data: Fig. 13.17 illustrates gender distribution with respect to organizer teams of Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In contrast, Fig. 13.18 shows the proportion of women in

Geschlechter-bezogene Daten: Fig. 13.17 enthält Daten zur Geschlechter-Verteilung in der Seminarleitung. Dagegen zeigt Fig. 13.18 die Quote von Frauen bei der Beantragung von Seminaren sowohl bezüglich der Teams als auch bezüglich der gesamten Antragsteller. Die Abbildungen Fig. 13.19 und Fig. 13.20 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. 13.21 aufgezeigt.

seminar proposals with respect to both the teams and the proposers overall. Fig. 13.19 and Fig. 13.20 mainly illustrate the share of female participants and invitees, respectively. The distribution of the share of women among seminar participants is displayed in Fig. 13.21.

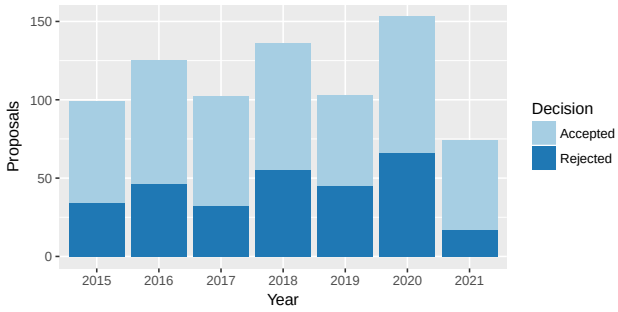

(a) Chart for 2015–2021

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
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
2020	153	87	56.9	66	43.1
2021	74	57	77.0	17	23.0

(b) Detailed numbers for 2015–2021

Fig. 13.1
Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

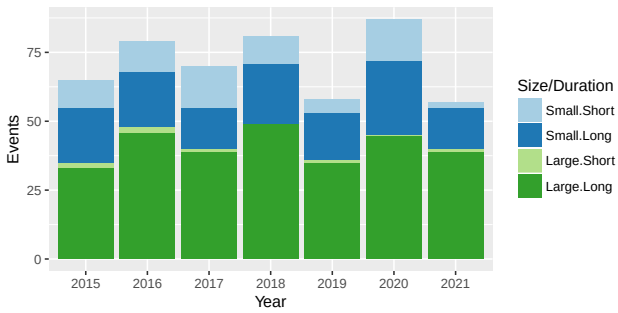

(a) Chart for 2015–2021

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
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
2020	15	27	0	45	87
2021	2	15	1	39	57

(b) Detailed numbers for 2015–2021

Fig. 13.2
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2015–2021. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

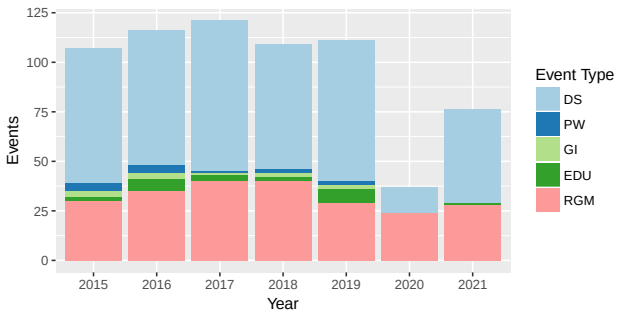

(a) Chart for 2015–2021

Year	DS	PW	GI	EDU	RGM	Total
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
2020	13	0	0	0	24	37
2021	47	0	0	1	28	76

(b) Detailed numbers for 2015–2021

Fig. 13.3
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.

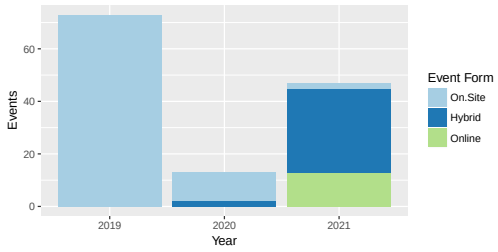

(a) Graphical distribution of seminars in group A in 2019–2021 by year and form

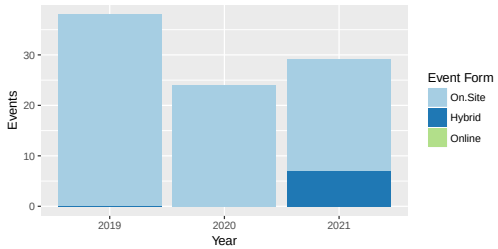

(b) Graphical distribution of seminars in group B in 2019–2021 by year and event form

	Group A			Group B		
Year	On-Site	Hybrid	Online	On-Site	Hybrid	Online
2019	73	0	0	38	0	0
2020	11	2	0	24	0	0
2021	2	32	13	22	7	0

(c) Detailed numbers for 2019–2021 by event form

Fig. 13.4
Number of all events held at Dagstuhl, by event form and group. Group A = Dagstuhl Seminars and Dagstuhl-Perspectives-Workshops, group B = all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).

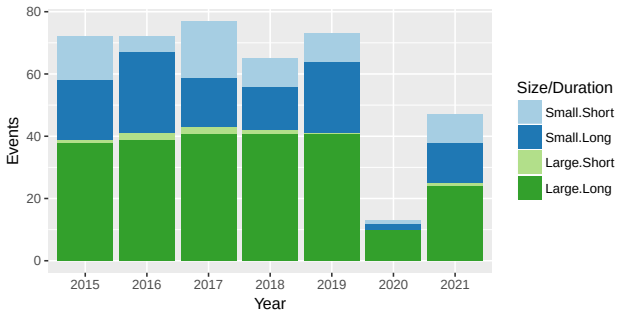

(a) Chart for 2015–2021

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
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
2020	1	2	0	10	13
2021	9	13	1	24	47

(b) Detailed numbers for 2015–2021

Fig. 13.5
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2015–2021. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

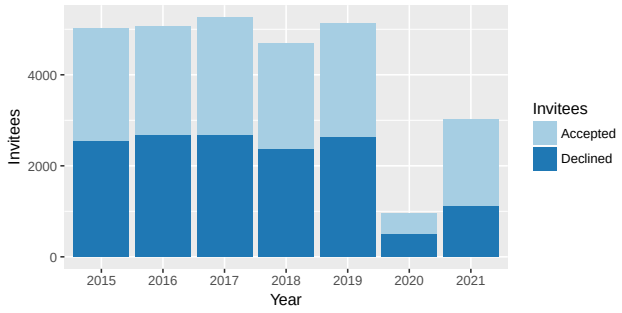

(a) Chart for 2015–2021

Year	Invitees	Accepted		Declined	
	#	#	%	#	%
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
2020	964	442	45.9	522	54.1
2021	3022	1894	62.7	1128	37.3

(b) Detailed numbers for 2015–2021

Fig. 13.6
Total number of invitees, accepted and declined invitations for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

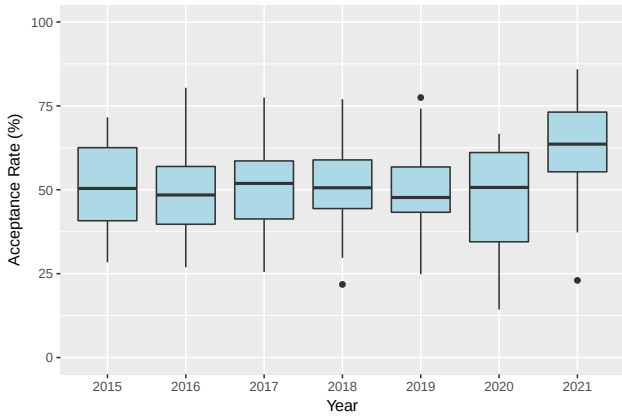

(a) Chart for 2015–2021

Year	Min (%)	Max (%)	Avg (%)	Std (%)
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
2020	14.3	66.7	45.8	17.4
2021	23.0	85.9	63.2	12.5

(b) Detailed numbers for 2015–2021

Fig. 13.7
Acceptance rate distribution per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2015–2021. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

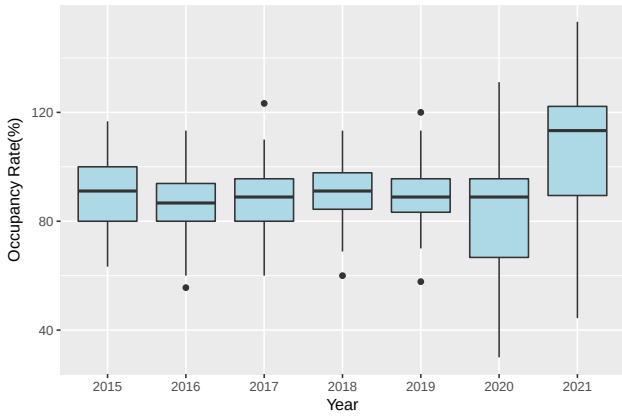

(a) Chart for 2015–2021

Year	Min (%)	Max (%)	Avg (%)	Std (%)
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
2020	30.0	131.1	79.7	27.9
2021	44.4	153.3	108.0	24.6

(b) Detailed numbers for 2015–2021

Fig. 13.8
Occupancy rate distribution per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2015–2021. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

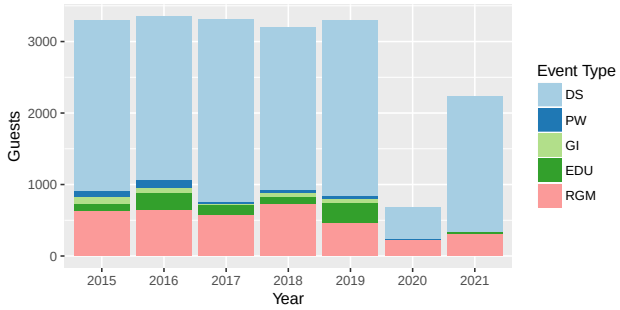

(a) Chart for 2015–2021

Year	DS		PW		GI		EDU		RGM		Total
	#	%	#	%	#	%	#	%	#	%	#
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
2020	442	65.1	0	0.0	0	0.0	0	0.0	237	34.9	679
2021	1894	84.6	0	0.0	0	0.0	31	1.4	314	14.0	2239

(b) Detailed numbers for 2015–2021

Fig. 13.9
Number of participants. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

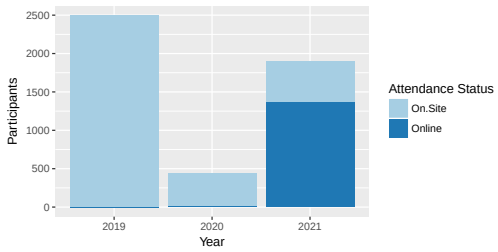

(a) Graphical distribution of participants of seminars in group A in 2019–2021 by year and attendance status

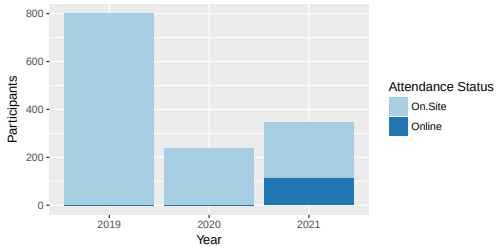

(b) Graphical distribution of participants of seminars in group B in 2019–2021 by year and attendance status

	Group A			Group B		
Year	On-Site	Online	Total	On-Site	Online	Total
2019	2498	0	2498	801	0	801
2020	419	23	442	237	0	237
2021	529	1365	1894	233	112	345

(c) Detailed numbers for 2019–2021 by attendance status

Fig. 13.10
Number of participants by attendance status and group. Group A = Dagstuhl Seminars and Dagstuhl-Perspectives-Workshops, group B = all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).

Country	A	B	Total
Germany	390	270	660
United States	389	9	398
United Kingdom	215	10	225
France	112	11	123
Netherlands	74	21	95
Italy	68	0	68
Switzerland	61	2	63
Austria	57	5	62
Canada	42	1	43
Sweden	43	0	43
Japan	42	0	42
Spain	34	2	36
Belgium	31	1	32
Israel	30	0	30
Poland	27	0	27
Australia	23	1	24
Czech Republic	24	0	24
Denmark	22	1	23
China	21	0	21
India	21	0	21
Norway	16	5	21
Finland	19	0	19
Greece	14	0	14
Hungary	11	0	11
Republic of Korea	11	0	11
Brazil	10	0	10

Country	A	B	Total
Singapore	10	0	10
Russian Federation	9	0	9
Slovenia	8	1	9
Ireland	8	0	8
Portugal	7	0	7
New Zealand	6	0	6
Cyprus	5	0	5
Bulgaria	4	0	4
Serbia	3	1	4
Argentina	3	0	3
Estonia	2	1	3
Hong Kong	3	0	3
Pakistan	0	3	3
Turkey	3	0	3
Chile	2	0	2
Latvia	2	0	2
Slovak Republic	2	0	2
Taiwan	2	0	2
Colombia	1	0	1
Ecuador	1	0	1
Lithuania	1	0	1
Luxembourg	1	0	1
Malta	1	0	1
Romania	1	0	1
Saudi Arabia	1	0	1
South Africa	1	0	1
Total	1894	345	2239

(a) Details for 2021 by country

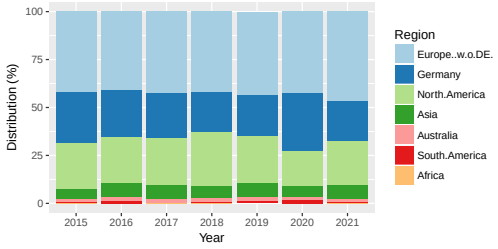

(b) Graphical distribution of seminars in group A in 2015–2021 by year and region

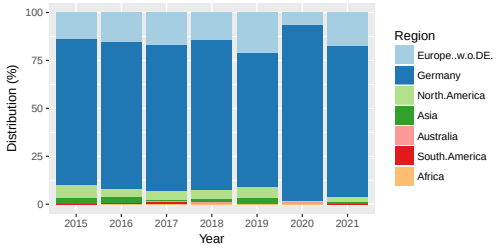

(c) Graphical distribution of seminars in group B in 2015–2021 by year and region

Region	Group A		Group B		Total	
	#	%	#	%	#	%
Europe (w/o Germany)	885	46.7	61	17.7	946	42.3
Germany	390	20.6	270	78.3	660	29.5
North America	431	22.8	10	2.9	441	19.7
Asia	141	7.4	3	0.9	144	6.4
Australia	29	1.5	1	0.3	30	1.3
South America	17	0.9	0	0	17	0.8
Africa	1	0.1	0	0	1	0
Total	1894	100	345	100	2239	100

(d) Details for 2021 by region

Fig. 13.11

Number of Dagstuhl guests by country of origin. Group A = Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, group B = Participants in all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).

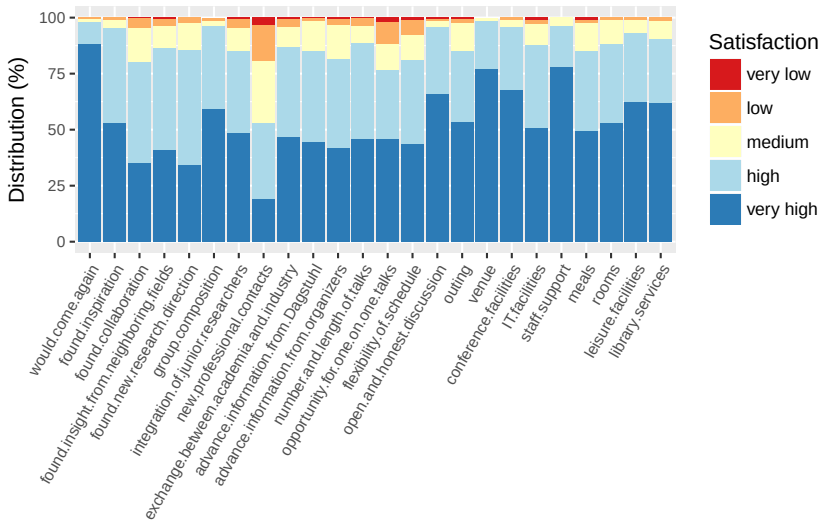

(a) Graphical distribution for 2021

	2015	2016	2017	2018	2019	2020	2021	2021 – Detailed Numbers					
	ø	ø	ø	ø	ø	ø	ø	1	2	3	4	5	total
would come again	4.9	4.9	4.9	4.9	4.9	4.9	4.9	0	3	8	65	578	654
found inspiration	4.5	4.5	4.5	4.6	4.6	4.6	4.5	0	7	22	278	347	654
found collaboration	4.1	4.1	4.2	4.2	4.3	4.3	4.1	1	27	98	283	225	634
found insight from neighboring fields	4.3	4.2	4.2	4.2	4.3	4.2	4.2	4	20	63	294	267	648
found new research direction	4.1	4.1	4.1	4.2	4.2	4.2	4.2	0	15	76	333	220	644
group composition	4.5	4.5	4.5	4.5	4.6	4.5	4.5	1	8	15	242	386	652
integration of junior researchers	4.2	4.3	4.3	4.3	4.3	4.3	4.3	4	23	68	233	313	641
new professional contacts	3.6	3.7	3.8	3.8	3.8	3.6	3.5	19	97	168	207	115	606
exchange between academia and industry	4.3	4.3	4.4	4.4	4.4	4.3	4.3	2	14	36	163	188	403
advance information from Dagstuhl	4.4	4.4	4.4	4.4	4.5	4.5	4.3	1	7	85	256	281	630
advance information from organizers	4.1	4.2	4.1	4.3	4.3	4.3	4.2	3	18	95	249	265	630
number and length of talks	4.2	4.3	4.2	4.3	4.3	4.3	4.3	2	21	50	280	300	653
opportunity for one on one talks	4.5	4.6	4.6	4.6	4.6	4.6	4.1	10	54	66	170	254	554
flexibility of schedule	4.3	4.4	4.3	4.3	4.5	4.4	4.2	6	39	69	228	266	608
open and honest discussion	4.7	4.7	4.7	4.7	4.7	4.7	4.6	3	5	19	191	422	640
outing	4.1	4.2	4.2	4.3	4.2	4.2	4.4	1	4	26	67	112	210
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.8	0	0	5	86	307	398
conference facilities	4.6	4.7	4.7	4.7	4.7	4.7	4.6	0	4	11	102	248	365
IT facilities	4.3	4.4	4.3	4.4	4.4	4.3	4.3	3	6	30	118	162	319
staff support	4.7	4.7	4.7	4.8	4.8	4.7	4.7	0	0	15	76	324	415
meals	4.1	4.1	4.1	4.1	4.2	4.1	4.3	3	4	39	112	155	313
rooms	4.4	4.4	4.4	4.4	4.4	4.3	4.4	0	3	34	109	167	313
leisure facilities	4.6	4.5	4.5	4.6	4.6	4.6	4.5	0	3	14	80	160	257
library services	4.5	4.5	4.5	4.6	4.6	4.5	4.5	0	2	11	40	86	139

(b) Averages for 2015–2021 and detailed numbers for 2021: 1 = very low, 2 = low, 3 = medium, 4 = high, 5 = very high

Fig. 13.12
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

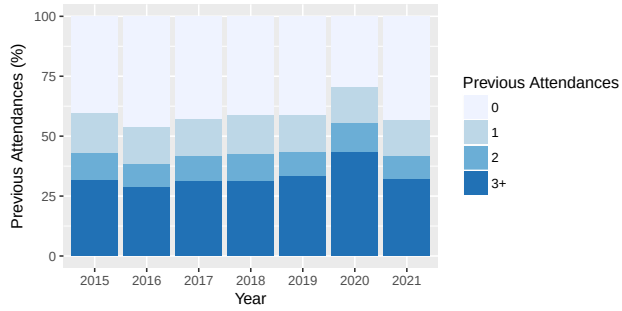

(a) Graphical distribution for 2015–2021

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	#
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
2020	61	29	31	15	25	12	90	43	207
2021	283	43	96	15	63	10	211	32	653

(b) Detailed numbers for 2015–2021

Fig. 13.13
Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants and their previous instances of attendance at Dagstuhl Seminars or Dagstuhl Perspectives Workshops, according to our guest survey.

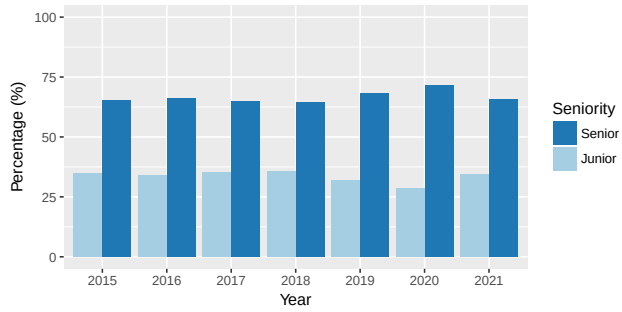

(a) Chart for 2015–2021

Year	Junior		Senior		Total
	#	%	#	%	#
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
2020	53	28.5	133	71.5	186
2021	195	34.5	370	65.5	565

(b) Detailed numbers for 2015–2021

Fig. 13.14
Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

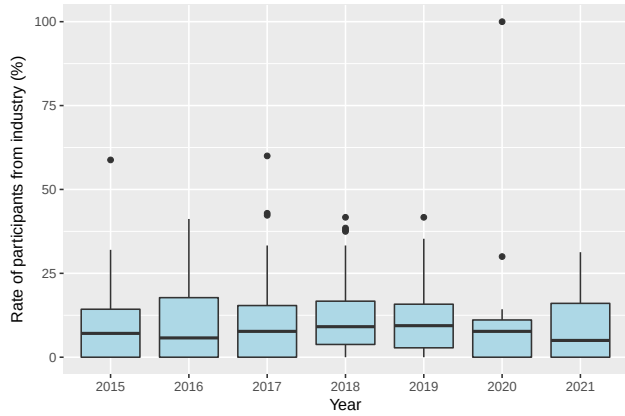

(a) Chart for 2015–2021

Year	Min (%)	Max (%)	Avg (%)	Std (%)
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
2020	0.0	100.0	14.6	25.9
2021	0.0	31.3	8.5	10.1

(b) Detailed numbers for 2015–2021

Fig. 13.15
Distribution of the ratio of participants with self-assigned primary occupation in business per Dagstuhl Seminar and Dagstuhl Perspectives Workshop in 2015–2021, 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".

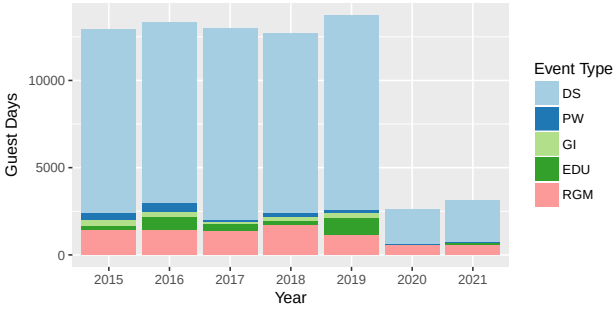

(a) Chart for 2015–2021

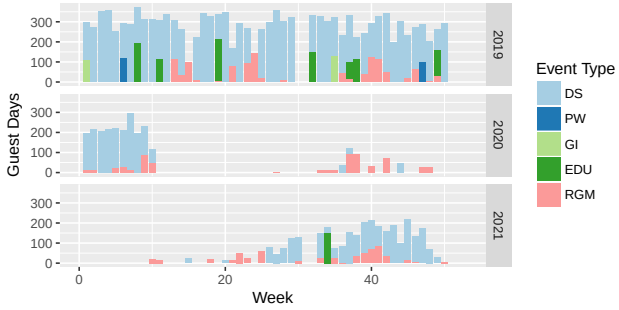

(b) Graphical distribution for 2019–2021 by week

Year	DS	PW	GI	EDU	RGM	Total
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
2020	1984	0	0	0	614	2598
2021	2397	0	0	150	576	3123

(c) Detailed numbers for 2015–2021

Fig. 13.16
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.

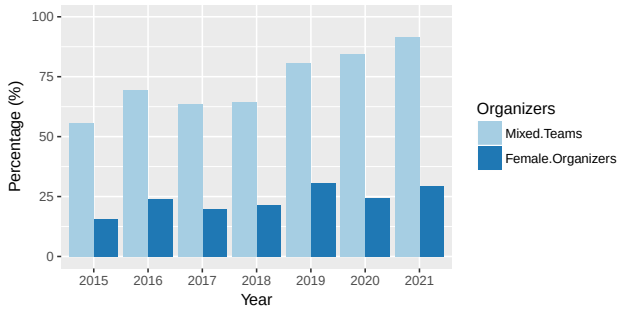

(a) Chart for 2015–2021

	Organizer Teams			Organizers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
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	42	64.6	259	56	21.6
2019	73	59	80.8	291	89	30.6
2020	13	11	84.6	53	13	24.5
2021	47	43	91.5	184	54	29.3

(b) Detailed numbers for 2015–2021

Fig. 13.17
Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

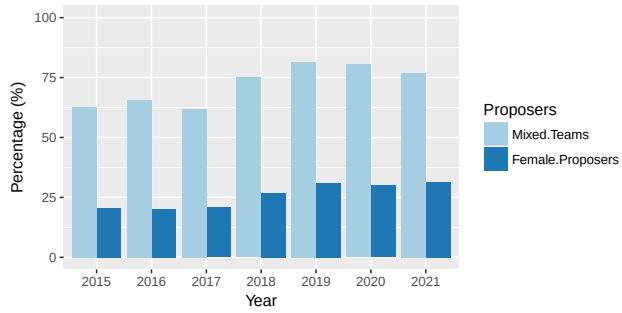

(a) Chart for 2015–2021

	Proposer Teams			Proposers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
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
2020	153	123	80.4	593	178	30.0
2021	74	57	77.0	296	93	31.4

(b) Detailed numbers for 2015–2021

Fig. 13.18
Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with mixed-gender proposer teams.

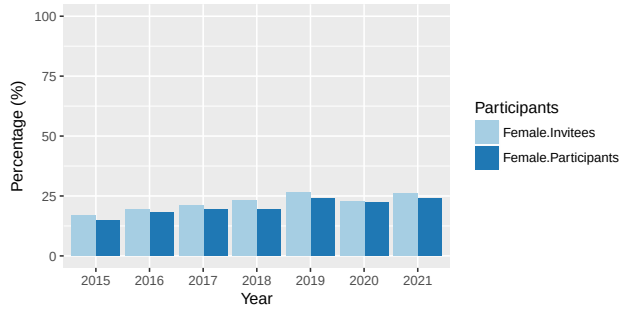

(a) Chart for 2015–2021

	Invitees			Participants		
Year	Total	Female		Total	Female	
	#	#	%	#	#	%
2015	5023	845	16.8	2473	369	14.9
2016	5060	977	19.3	2393	437	18.3
2017	5267	1110	21.1	2572	495	19.2
2018	4692	1086	23.1	2320	453	19.5
2019	5143	1368	26.6	2498	603	24.1
2020	964	218	22.6	442	99	22.4
2021	3022	793	26.2	1894	455	24.0

(b) Detailed numbers for 2015–2021

Fig. 13.19
Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.

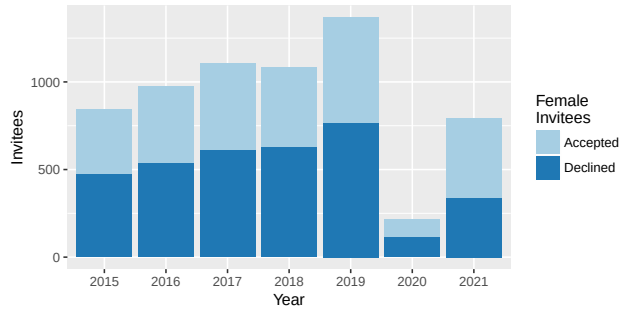

(a) Chart for 2015–2021

Year	Female Invitees	Accepted		Declined	
	#	#	%	#	%
2015	845	369	43.7	476	56.3
2016	977	437	44.7	540	55.3
2017	1110	495	44.6	615	55.4
2018	1086	453	41.7	633	58.3
2019	1368	603	44.1	765	55.9
2020	218	99	45.4	119	54.6
2021	793	455	57.4	338	42.6

(b) Detailed numbers for 2015–2021

Fig. 13.20
Female invitees to Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

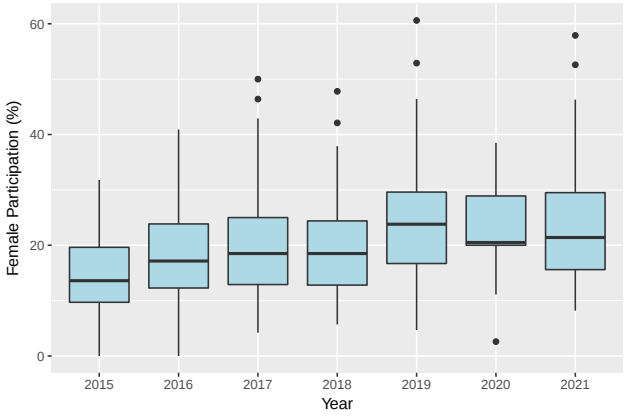

(a) Chart for 2015–2021

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2015	0.0	31.8	14.8	7.7
2016	0.0	40.9	18.3	9.1
2017	4.2	50.0	19.7	9.8
2018	5.7	47.8	19.8	9.2
2019	4.7	60.6	24.6	10.6
2020	2.6	38.5	22.7	9.2
2021	8.2	57.9	24.5	11.4

(b) Detailed numbers for 2015–2021

Fig. 13.21
Distribution of the share of women among participants per Dagstuhl Seminar or Dagstuhl Perspectives Workshop in 2015–2021. 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. 13.22 listet die durchschnittlichen Nutzungszahlen der letzten Jahre auf. Ein Überblick über die Entwicklung des dblp Datenbestandes kann Fig. 13.23 und Fig. 13.24 entnommen werden. Fig. 13.25–13.27 geben Auskunft über die kontinuierliche Datenkuration und -anreicherung des Bestandes.

This section provides statistical data about the dblp computer science bibliography. Fig. 13.22 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. 13.23 and Fig. 13.24. Information about the continuous data curation and enrichment of existing records can be found in Fig. 13.25–13.27.

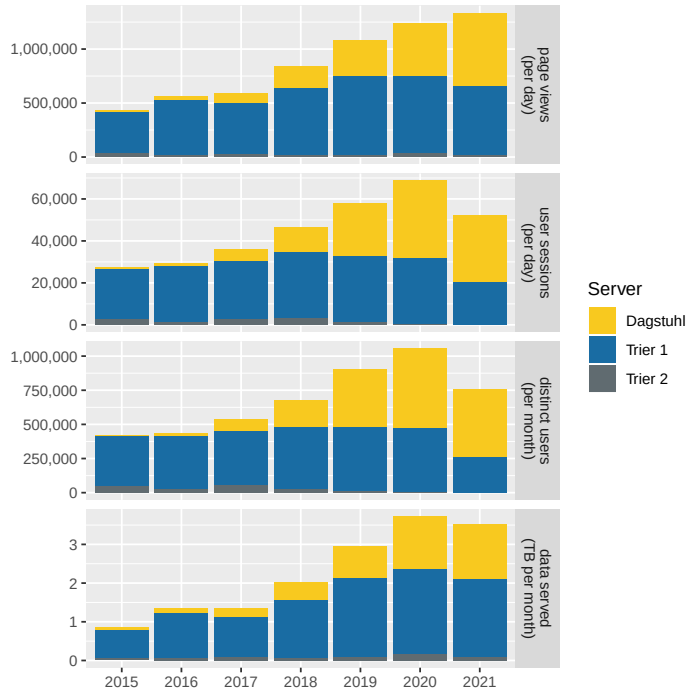

(a) Chart for 2015–2021

	Trier 1		Trier 2		Dagstuhl		Total		
	2020	2021	2020	2021	2020	2021	2020	2021	%
page views per day	712,228	644,837	40,523	18,206	486,166	672,274	1,238,918	1,335,318	+7.8
user sessions (visits) per day	31,450	20,134	414	297	36,849	31,537	68,715	51,969	-24.4
page views per user session	22.6	32.0	97.7	61.1	13.2	21.3	18.0	25.7	+42.5
distinct users (IPs) per month	473,014	257,769	4,433	3,512	578,492	492,268	1,055,941	753,550	-28.6
data served per month	2,243.6 GB	2,073.9 GB	170.3 GB	88.4 GB	1,402.1 GB	1,433.5 GB	3,816.0 GB	3,595.8 GB	-5.8

(b) Detailed numbers for the past two years

Fig. 13.22
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. Since 2017, server Dagstuhl has been promoted to play a more prominent role under the domain dblp.org. In 2021, the counting method to determine distinct users has been changed in order to avoid double counting of the same IP across diferent servers.

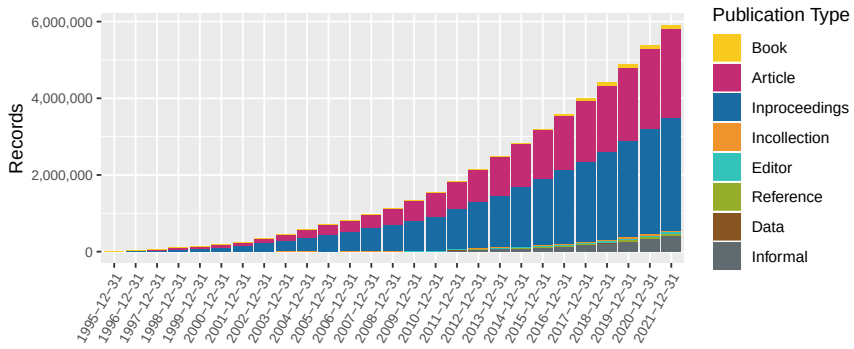

(a) Chart for 1996–2021

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
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
2020	95,442	1.8	2,096,005	38.9	2,735,716	50.8	36,372	0.7	48,481	0.9	27,321	0.5	2,230	0.0	345,676	6.4	5,387,243
2021	99,289	1.7	2,310,800	39.1	2,951,745	50.0	37,745	0.6	51,968	0.9	27,321	0.5	2,882	0.0	425,791	7.2	5,907,541

(b) Detailed numbers for 2015–2021

Fig. 13.23
Development of the total size of the dblp database.

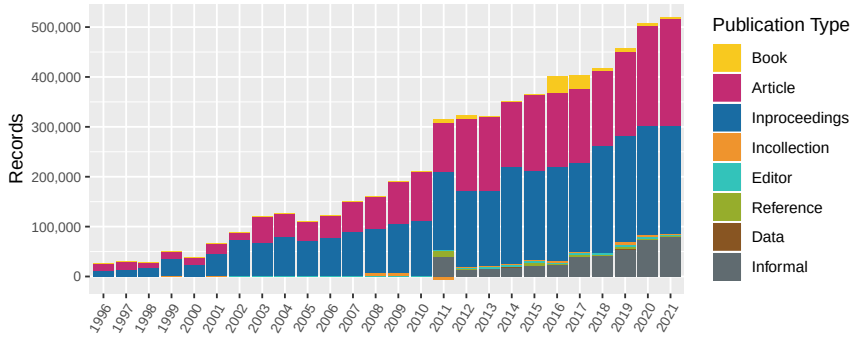

(a) Chart for 1996–2021

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Data		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
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
2020	4,616	0.9	199,991	39.5	217,418	42.9	5,915	1.2	3,583	0.7	324	0.1	828	0.2	74,043	14.6	506,718
2021	3,847	0.7	214,795	41.3	216,029	41.5	1,373	0.3	3,487	0.7	0	0.0	652	0.1	80,115	15.4	520,298

(b) Detailed numbers for 2015–2021

Fig. 13.24
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*.

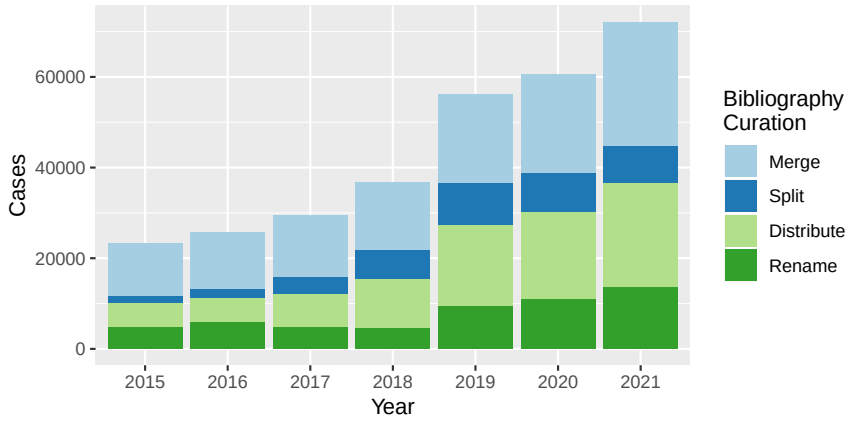

(a) Chart for 2015–2021

Year	Merge		Split		Distribute		Rename		Total
	#	%	#	%	#	%	#	%	#
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	19,595	34.9	9,192	16.4	17,795	31.7	9,562	17.0	56,144
2020	21,636	35.7	8,606	14.2	19,326	31.9	11,083	18.3	60,651
2021	27,204	37.7	8,254	11.4	22,877	31.7	13,785	19.1	72,120

(b) Detailed numbers for 2015–2021

Fig. 13.25

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.

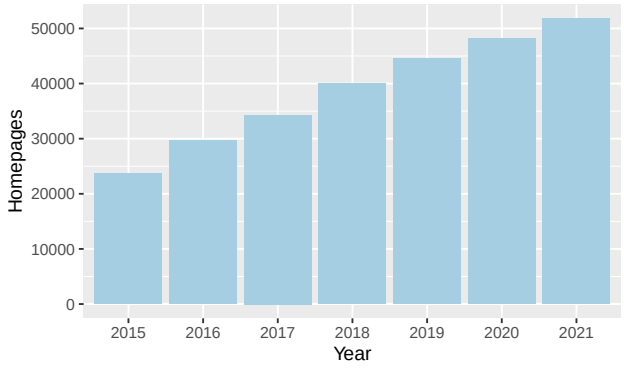

(a) Chart for 2015–2021

Year	Homepages
2015	23,698
2016	29,642
2017	34,295
2018	40,046
2019	44,562
2020	48,216
2021	51,756

(b) Detailed numbers for 2015–2021

Fig. 13.26

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 13.27.

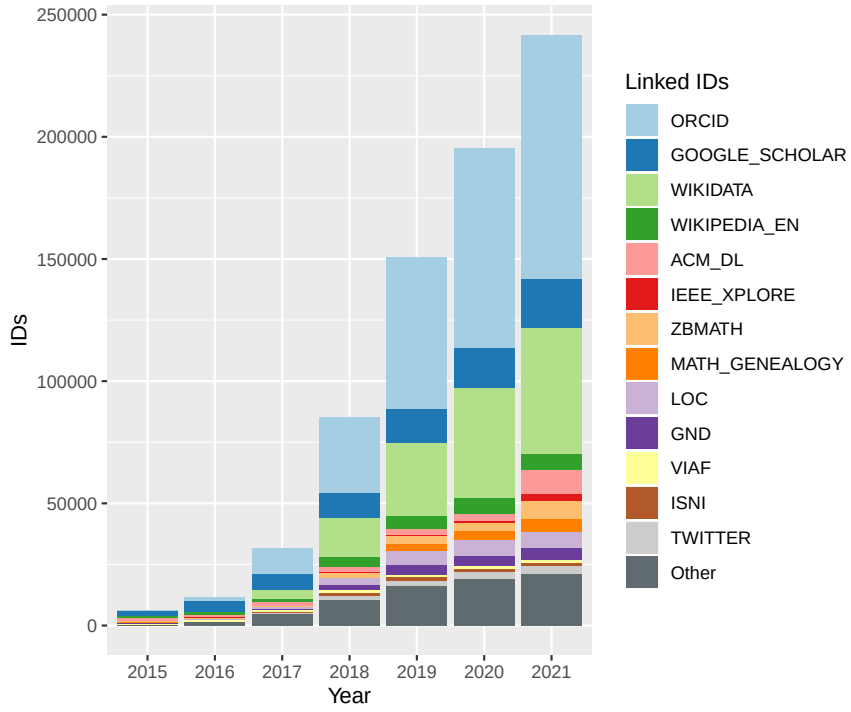

(a) Chart for 2015–2021

Year	ORCID		Google Scholar		Wikidata		Wikipedia (en)		ACM DL		IEEE Xplore		zbMATH	
	#	%	#	%	#	%	#	%	#	%	#	%	#	%
2015	73	1.2	2,125	34.2	0	0.0	902	14.5	1,095	17.6	0	0.0	568	9.2
2016	1,420	12.2	4,522	38.9	4	0.0	1,103	9.5	1,107	9.5	0	0.0	633	5.5
2017	10,342	32.8	6,608	21.0	3,588	11.4	1,376	4.4	1,277	4.1	0	0.0	646	2.0
2018	30,939	36.4	10,182	12.0	15,565	18.3	4,366	5.1	2,144	2.5	0	0.0	2,145	2.5
2019	61,976	41.1	13,726	9.1	30,022	19.9	5,547	3.7	2,448	1.6	104	0.1	3,363	2.2
2020	81,675	41.8	16,338	8.4	45,216	23.1	6,275	3.2	3,166	1.6	557	0.3	3,607	1.8
2021	99,459	41.2	20,060	8.3	51,744	21.4	6,668	2.8	9,642	4.0	3,013	1.2	7,486	3.1

Year	Math Genealogy		LOC		GND		VIAF		ISNI		Twitter		Other		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2015	7	0.1	324	5.2	217	3.5	516	8.3	1	0.0	62	1.0	315	5.1	6,205
2016	84	0.7	324	2.8	220	1.9	526	4.5	7	0.1	161	1.4	1,503	12.9	11,614
2017	5	0.0	602	1.9	430	1.4	840	2.7	606	1.9	392	1.2	4,804	15.2	31,516
2018	65	0.1	2,862	3.4	2,385	2.8	938	1.1	1,231	1.4	1,567	1.8	10,670	12.5	85,059
2019	3,071	2.0	5,680	3.8	3,782	2.5	1,060	0.7	1,383	0.9	2,183	1.4	16,349	10.8	150,694
2020	3,590	1.8	6,289	3.2	4,241	2.2	1,061	0.5	1,385	0.7	2,613	1.3	19,358	9.9	195,371
2021	5,048	2.1	6,877	2.8	4,777	2.0	1,068	0.4	1,390	0.6	3,183	1.3	21,233	8.8	241,648

(b) Detailed numbers for 2015–2021

Fig. 13.27
Linked and verified external person IDs in dblp author bibliographies. A single bibliography may be linked to multiple external IDs.

Statistiken zu Dagstuhl Publishing

13.3 Statistics of Dagstuhl Publishing

Dieser Abschnitt enthält statistische Daten zum Publikationswesen von Schloss Dagstuhl.

Ein Überblick über die Entwicklung der seminarbezogenen Veröffentlichungen kann den ersten drei Diagrammen und Tabellen entnommen werden. Fig. 13.28 fasst die statistischen Daten der Veröffentlichungen in der Zeitschrift Dagstuhl Reports zusammen, Fig. 13.29 die der Publikationen in der Reihe Dagstuhl Manifestos und schließlich Fig. 13.30 die der veröffentlichten Bände in der Reihe Dagstuhl Follow-Ups.

Die statistischen Daten zu den dienstleistungsbezogenen Veröffentlichungen finden sich anschließend: Fig. 13.31 fasst die Daten in der Reihe OASIcs und Fig. 13.32 die der Reihe LIPIcs zusammen.

Die Kennzahlen der Zeitschrift LITES können Fig. 13.33 entnommen werden.

Die verschiedenen Publikationsserien wurden in unterschiedlichen Jahren zwischen 2009 und 2015 gegründet. Wir stellen in den Statistiken dennoch stets den gesamten Zeitraum (2015–2021) dar.

In this section the statistical data of Dagstuhl Publishing is presented.

The first three figures present the development of the seminar-focused series: Fig. 13.28 summarizes the data of the periodical Dagstuhl Reports, Fig. 13.29 the data of the Dagstuhl Manifestos series, and, finally, Fig. 13.30 that of the volumes published in the Dagstuhl Follow-Ups series.

The statistical data of the service-focused series are presented afterwards. Fig. 13.31 presents numbers related to OASIcs and Fig. 13.32 numbers related to LIPIcs.

We summarize the publications of the journal LITES in Fig. 13.33.

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 (2015–2021).

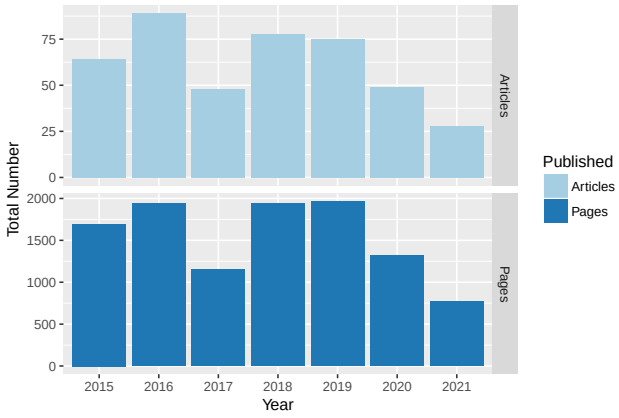

(a) Graphical distribution for 2015–2021

Year	Articles	Pages
2015	64	1695
2016	89	1940
2017	48	1154
2018	78	1938
2019	75	1959
2020	49	1322
2021	28	770

(b) Detailed numbers for 2015–2021

Fig. 13.28
Statistics about Dagstuhl Reports published between 2015 to 2021.

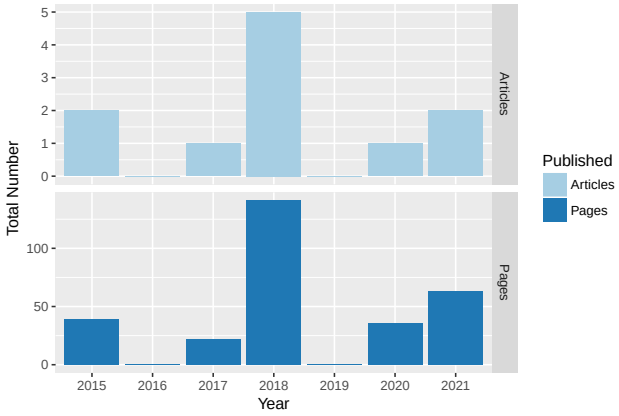

(a) Graphical distribution for 2015–2021

Year	Articles	Pages
2015	2	39
2016	0	0
2017	1	22
2018	5	141
2019	0	0
2020	1	36
2021	2	63

(b) Detailed numbers for 2015–2021

Fig. 13.29
Statistics about Dagstuhl Manifestos published between 2015 to 2021.

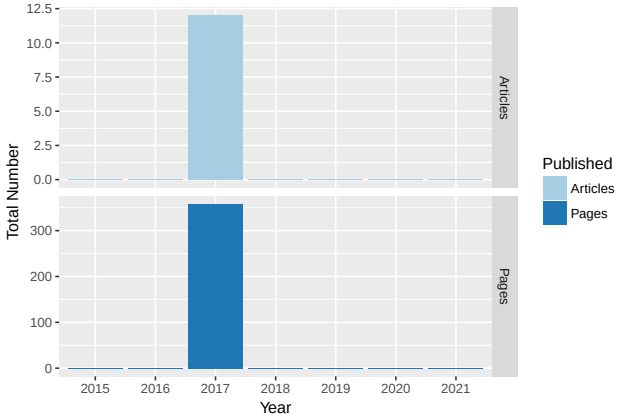

(a) Graphical distribution for 2015–2021

Year	Volumes	Articles	Pages
2015	0	0	0
2016	0	0	0
2017	1	12	358
2018	0	0	0
2019	0	0	0
2020	0	0	0
2021	0	0	0

(b) Detailed numbers for 2015–2021

Fig. 13.30
Statistics about Dagstuhl Follow-Ups volumes published between 2015 to 2021.

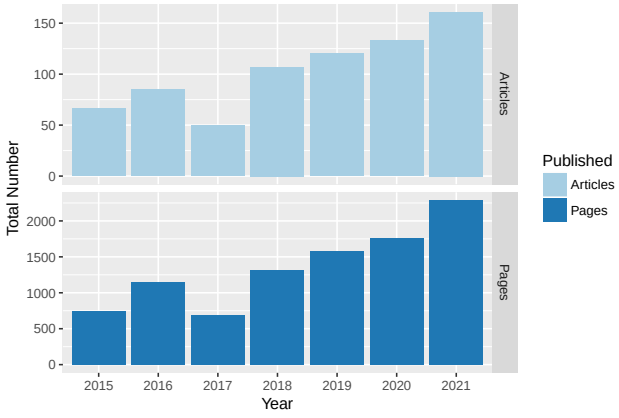

(a) Graphical distribution for 2015–2021

Year	Volumes	Articles	Pages
2015	6	66	750
2016	6	85	1152
2017	3	50	694
2018	7	107	1312
2019	9	120	1576
2020	11	133	1754
2021	11	161	2284

(b) Detailed numbers for 2015–2021

Fig. 13.31
Statistics about OASIcs volumes published between 2015 to 2021.

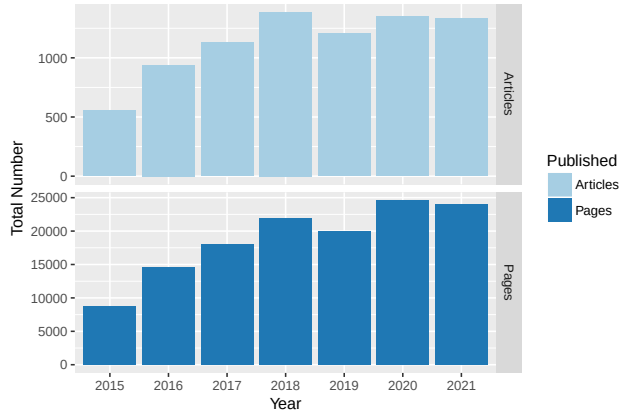

(a) Graphical distribution for 2015–2021

Year	Volumes	Articles	Pages
2015	16	555	8805
2016	19	939	14538
2017	25	1127	18082
2018	32	1387	21876
2019	29	1208	20032
2020	32	1352	24562
2021	32	1333	23960

(b) Detailed numbers for 2015–2021

Fig. 13.32
Statistics about LIPIcs volumes published between 2015 to 2021.

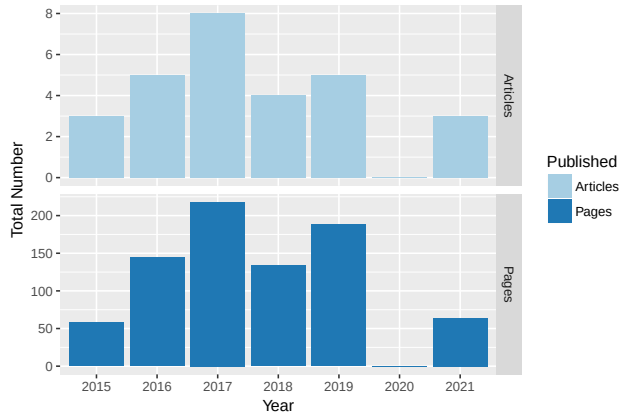

(a) Graphical distribution for 2015–2021

Year	Articles	Pages
2015	3	58
2016	5	144
2017	8	218
2018	4	134
2019	5	188
2020	0	0
2021	3	64

(b) Detailed numbers for 2015–2021

Fig. 13.33
Statistics about LITES articles published between 2015 to 2021.

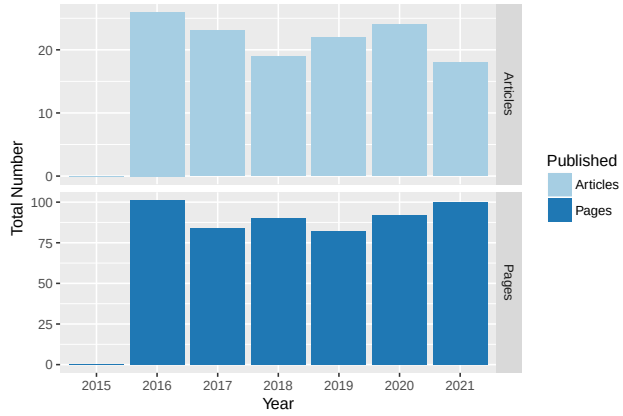

(a) Graphical distribution for 2015–2021

Year	Articles	Pages
2015	0	0
2016	26	101
2017	23	84
2018	19	90
2019	22	82
2020	24	92
2021	18	100

(b) Detailed numbers for 2015–2021

Fig. 13.34
Statistics about DARTS artifacts published between 2015 to 2021.

14 Veranstaltungen 2021
Schedule of Events 2021

Dagstuhl-Seminare 14.1 *Dagstuhl Seminars*

21071 – Scalable Data Structures
Gerth Stølting Brodal (Aarhus University, DK), John Iacono (UL – Brussels, BE), Markus E. Nebel (Universität Bielefeld, DE), Vijaya Ramachandran (University of Texas – Austin, US)
February 14–19, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21071

21072 – Autonomous Agents on the Web
Olivier Boissier (Ecole des Mines – St. Etienne, FR), Andrei Ciortea (Universität St. Gallen, CH), Andreas Harth (Fraunhofer IIS – Nürnberg, DE), Alessandro Ricci (Università di Bologna, IT)
February 14–19, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21072

21121 – Computational Complexity of Discrete Problems
Anna Gál (University of Texas – Austin, US), Meena Mahajan (Institute of Mathematical Sciences – Chennai, IN), Rahul Santhanam (University of Oxford, GB), Till Tantau (Universität zu Lübeck, DE)
March 21–26, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21121

21152 – Multi-Level Graph Representation for Big Data Arising in Science Mapping
Katy Börner (Indiana University – Bloomington, US), Stephen G. Kobourov (University of Arizona – Tucson, US)
April 11–16, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21152

21171 – Temporal Graphs: Structure, Algorithms, Applications
Arnaud Casteigts (University of Bordeaux, FR), Kitty Meeks (University of Glasgow, GB), George B. Mertzios (Durham University, GB), Rolf Niedermeier (TU Berlin, DE)
April 25–30, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21171

21181 – Computational Geometry
Siu-Wing Cheng (HKUST – Kowloon, HK), Anne Driemel (Universität Bonn, DE), Jeff M. Phillips (University of Utah – Salt Lake City, US)
May 2–7, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21181

21192 – Approaches and Applications of Inductive Programming
Andrew Cropper (University of Oxford, GB), Luc De Raedt (KU Leuven, BE), Richard Evans (DeepMind – London, GB), Ute Schmid (Universität Bamberg, DE)
May 9–12, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21192

21201 – Serverless Computing
Cristina Abad (ESPOL – Guayaquil, EC), Ian T. Foster (Argonne National Laboratory – Lemont, US), Nikolas Herbst (Universität Würzburg, DE), Alexandru Iosup (VU University Amsterdam, NL)
May 16–21, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21201

21231 – Transparency by Design
Casey Dugan (IBM Research – Cambridge, US), Judy Kay (The University of Sydney, AU), Tsvi Kuflik (Haifa University, IL), Michael Rovatsos (University of Edinburgh, GB)
June 6–10, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21231

21232 – Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments
Stephen Brewster (University of Glasgow, GB), Andrew Kun (University of New Hampshire – Durham, US), Andreas Riener (TH Ingolstadt, DE), Orit Shaer (Wellesley College, US)
June 6–11, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21232

21243 – Compute-First Networking
Jon Crowcroft (University of Cambridge, GB), Philip Eardley (BT Applied Research – Ipswich, GB), Dirk Kutscher (FH Emden, DE), Eve M. Schooler (Intel – Santa Clara, US)
June 13–16, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21243

21261 – Quantum Complexity: Theory and Application
Bill Fefferman (University of Chicago, US), Sevag Gharibian (Universität Paderborn, DE), Norbert Schuch (Universität Wien, AT), Barbara Terhal (TU Delft, NL)
June 27 to July 2, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21261

21262 – Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users
Ana Aguiar (Universidade do Porto, PT), Onur Altintas (Toyota Motors North America – Mountain View, US), Falko Dressler (TU Berlin, DE), Gunnar Karlsson (KTH Royal Institute of Technology – Kista, SE)
June 28, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21262

21271 – Computational Proteomics
Sebastian Böcker (Universität Jena, DE), Rebekah Gundry (University of Nebraska – Omaha, US), Lennart Martens (Ghent University, BE), Magnus Palmblad (Leiden University Medical Center, NL)
July 4–9, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21271

21272 – Towards Climate-Friendly Internet Research
Vaibhav Bajpai (TU München, DE), Jon Crowcroft (University of Cambridge, GB), Oliver Hohlfeld (BTU Cottbus, DE), Srinivasan Keshav (University of Cambridge, GB)
July 6–9, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21272

21283 – Data Structures for Modern Memory and Storage Hierarchies
Stratos Idreos (Harvard University – Cambridge, US), Viktor Leis (Universität Erlangen-Nürnberg, DE), Kai-Uwe Sattler (TU Ilmenau, DE), Margo Seltzer (University of British Columbia – Vancouver, CA)
July 11–16, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21283

21292 – Scalable Handling of Effects
Danel Ahman (University of Ljubljana, SI), Amal Ahmed (Northeastern University – Boston, US), Sam Lindley (University of Edinburgh, GB), Andreas Rossberg (Dfinity – Zürich, CH)
July 18–23, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21292

21293 – Parameterized Complexity in Graph Drawing
Robert Ganian (TU Wien, AT), Fabrizio Montecchiani (University of Perugia, IT), Martin Nöllenburg (TU Wien, AT), Meirav Zehavi (Ben Gurion University – Beer Sheva, IL)
July 18–23, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21293

21301 – Matching Under Preferences: Theory and Practice
Haris Aziz (UNSW – Sydney, AU), Péter Biró (Hungarian Academy of Sciences – Budapest, HU), Tamás Fleiner (Budapest University of Technology & Economics, HU), Bettina Klaus (University of Lausanne, CH), David Manlove (University of Glasgow, GB)
July 25–30, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21301

21302 – Approximate Systems
Eva Darulova (MPI-SWS – Kaiserslautern, DE), Babak Falsafi (EPFL – Lausanne, CH), Andreas Gerstlauer (University of Texas at Austin, US), Phillip Stanley-Marbell (University of Cambridge, GB)
July 25–30, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21302

21331 – Coalition Formation Games
Edith Elkind (University of Oxford, GB), Judy Goldsmith (University of Kentucky – Lexington, US), Anja Rey (Universität Köln, DE), Jörg Rothe (Heinrich-Heine-Universität Düsseldorf, DE)
August 15–20, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21331

21332 – Understanding I/O Behavior in Scientific and Data-Intensive Computing
Philip Carns (Argonne National Laboratory, US), Julian Kunkel (Gesellschaft f. wissenschaftl. Datenverarbeitung, DE), Kathryn Mohror (LLNL – Livermore, US), Martin Schulz (TU München, DE)
August 15–20, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21332

21342 – Identifying Key Enablers in Edge Intelligence
Aaron Ding (TU Delft, NL), Ella Peltonen (University of Oulu, FI), Sasu Tarkoma (University of Helsinki, FI), Lars Wolf (TU Braunschweig, DE)
August 22–25, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21342

21351 – Universals of Linguistic Idiosyncrasy in Multilingual Computational Linguistics
Timothy Baldwin (The University of Melbourne, AU), William Croft (University of New Mexico – Alburquerque, US), Joakim Nivre (Uppsala University, SE), Agata Savary (Université de Tours – Blois, FR)
August 30–31, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21351

21352 – Higher-Order Graph Models: From Theoretical Foundations to Machine Learning
Tina Eliassi-Rad (Northeastern University – Boston, US), Vito Latora (Queen Mary University of London, GB), Martin Rosvall (University of Umeå, SE), Ingo Scholtes (Universität Würzburg, DE & Universität Zürich, CH)
August 29 to September 1, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21352

21361 – Extending the Synergies Between SAT and Description Logics
Joao Marques-Silva (CNRS – Toulouse, FR), Rafael Penaloza (University of Milano-Bicocca, IT), Uli Sattler (University of Manchester, GB)
September 5–10, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21361

21362 – Structure and Learning
Tiansi Dong (Universität Bonn, DE), Achim Rettinger (Universität Trier, DE), Jie Tang (Tsinghua University – Beijing, CN), Barbara Tversky (Columbia University – New York, US), Frank van Harmelen (VU University Amsterdam, NL)
September 5–10, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21362

21371 – Integrated Deduction
Maria Paola Bonacina (University of Verona, IT), Claudia Nalon (University of Brasilia, BR), Philipp Rümmer (Uppsala University, SE), Renate Schmidt (University of Manchester, GB)
September 12–17, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21371

21372 – Behavioural Types: Bridging Theory and Practice
Mariangiola Dezani (University of Turin, IT), Roland Kuhn (Actyx AG – München, DE), Sam Lindley (University of Edinburgh, GB), Alceste Scalas (Technical University of Denmark – Lyngby, DK)
September 12–17, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21372

21381 – Conversational Agent as Trustworthy Autonomous System (Trust-CA)
Asbjorn Folstad (SINTEF – Oslo, NO), Jonathan Grudin (Microsoft – Redmond, US), Effie Lai-Chong Law (Durham University, GB), Björn Schuller (Universität Augsburg, DE)
September 19–24, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21381

21391 – Sparsity in Algorithms, Combinatorics, and Logic
Daniel Král' (Masaryk University – Brno, CZ), Michal Pilipczuk (University of Warsaw, PL), Sebastian Siebertz (Universität Bremen, DE), Blair D. Sullivan (University of Utah – Salt Lake City, US)
September 26 to October 1, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21391

21401 – Visualization of Biological Data – From Analysis to Communication
Karsten Klein (Universität Konstanz, DE), Georgeta Elisabeta Marai (University of Illinois – Chicago, US), Kay Katja Nieselt (Universität Tübingen, DE), Blaz Zupan (University of Ljubljana, SI)
October 3–8, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21401

21402 – Digital Disinformation: Taxonomy, Impact, Mitigation, and Regulation
Claude Kirchner (INRIA – Le Chesnay, FR), Ninja Marnau (CISPA – Saarbrücken, DE), Franziska Roesner (University of Washington – Seattle, US)
October 3–6, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21402

21411 – Machine Learning in Sports
Ulf Brefeld (Universität Lüneburg, DE), Jesse Davis (KU Leuven, BE), Martin Lames (TU München, DE), Jim Little (University of British Columbia – Vancouver, CA)
October 10–15, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21411

21421 – Quantum Cryptanalysis
Stacey Jeffery (CWI – Amsterdam, NL), Michele Mosca (University of Waterloo, CA), Maria Naya-Plasencia (INRIA – Paris, FR), Rainer Steinwandt (University of Alabama in Huntsville, US)
October 17–22, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21421

21431 – Rigorous Methods for Smart Contracts
Nikolaj S. Bjørner (Microsoft – Redmond, US), Maria Christakis (MPI-SWS – Kaiserslautern, DE), Matteo Maffei (TU Wien, AT), Grigore Rosu (University of Illinois – Urbana-Champaign, US)
October 24–29, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21431

21432 – Probabilistic Numerical Methods – From Theory to Implementation
Philipp Hennig (Universität Tübingen, DE), Ilse Ipsen (North Carolina State University – Raleigh, US), Maren Mahsereci (Universität Tübingen, DE), Tim Sullivan (University of Warwick – Coventry, GB)
October 24–29, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21432

21441 – Adaptive Resource Management for HPC Systems
Michael Gerndt (TU München, DE), Masaaki Kondo (Keio University – Yokohama, JP), Barton P. Miller (University of Wisconsin-Madison, US), Tapasya Patki (LLNL – Livermore, US)
November 1–5, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21441

21442 – Ensuring the Reliability and Robustness of Database Management Systems
Alexander Böhm (SAP SE – Walldorf, DE), Maria Christakis (MPI-SWS – Kaiserslautern, DE), Eric Lo (The Chinese University of Hong Kong, HK), Manuel Rigger (ETH Zürich, CH)
November 1–4, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21442

21451 – Managing Industrial Control Systems Security Risks for Cyber Insurance
Simon Dejung (SCOR – Zürich, CH), Mingyan Liu (University of Michigan – Ann Arbor, US), Arndt Lüder (Otto-von-Guericke-Universität Magdeburg, DE), Edgar Weippl (University of Vienna & SBA Research – Wien)
November 7–12, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21451

21452 – Unambiguity in Automata Theory
Thomas Colcombet (CNRS – Paris, FR), Karin Quaas (Universität Leipzig, DE), Michal Skrzypczak (University of Warsaw, PL)
November 7–12, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21452

21461 – Descriptive Set Theory and Computable Topology
Mathieu Hoyrup (LORIA & INRIA Nancy, FR), Arno Pauly (Swansea University, GB), Victor Selivanov (A. P. Ershov Institute – Novosibirsk, RU), Mariya I. Soskova (University of Wisconsin – Madison, US)
November 14–19, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21461

21462 – Foundations of Persistent Programming
Hans-J. Boehm (Google – Mountain View, US), Ori Lahav (Tel Aviv University, IL), Azalea Raad (Imperial College London, GB)
November 14–17, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21462

21471 – Geometric Modeling: Interoperability and New Challenges
Falai Chen (Univ. of Science & Technology of China – Anhui, CN), Tor Dokken (SINTEF – Oslo, NO), Géraldine Morin (IRIT – University of Toulouse, FR), Tim Strotman (Siemens – Milford, US)
November 21–26, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21471

21472 – Geometric Logic, Constructivisation, and Automated Theorem Proving
Thierry Coquand (University of Gothenburg, SE), Hajime Ishihara (JAIST – Ishikawa, JP), Sara Negri (University of Genova, IT), Peter M. Schuster (University of Verona, IT)
November 21–26, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21472

21481 – Secure Compilation
David Chisnall (Microsoft Research – Cambridge, GB), Deepak Garg (MPI-SWS – Saarbrücken, DE), Catalin Hritcu (MPI-SP – Bochum, DE), Mathias Payer (EPFL – Lausanne, CH)
November 28 to December 3, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21481

21492 – Representing and Solving Spatial Problems
Pedro Cabalar (University of Coruña, ES), Zoe Falomir (Universitat Jaume I – Castello de la Plana, ES), Christian Freksa (in memoriam; † November 12, 2020), Paulo E. Santos (Flinders University – Adelaide, AU), Thora Tenbrink (Bangor University, GB)
December 5–10, 2021 | Dagstuhl Seminar | https://www.dagstuhl.de/21492

Dagstuhl-Perspektiven-Workshops 14.2 *Dagstuhl Perspectives Workshops*

There were no such meetings in 2021.

GI-Dagstuhl-Seminare 14.3 *GI-Dagstuhl Seminars*

There were no such meetings in 2021.

Lehrveranstaltungen 14.4 *Educational Events*

21343 – Leibniz MMS Summer School
Christian Bayer (Weierstraß Institut – Berlin, DE), Martin Eigel (Weierstraß Institut – Berlin, DE), Nicole Mücke (TU Braunschweig, DE), Feliks Nüske (Universität Paderborn, DE), David Sommer (Weierstraß Institut – Berlin, DE), Nikolas Tapia (Weierstraß Institut – Berlin, DE), Jia-Jie Zhu (Weierstraß Institut – Berlin, DE)
August 22–27, 2021 | Educational Event | https://www.dagstuhl.de/21343

Forschungsgruppentreffen 14.5 *Research Group Meetings*

21103 – Explainable Artificial Intelligence
Bernhard Ganter (TU Dresden, DE), Anatol Reibold (von Vicht GmbH, DE), Karl Erich Wolff (Ernst-Schröder-Zentrum Darmstadt, DE)
March 10–12, 2021 | Research Group Meeting | https://www.dagstuhl.de/21103

21118 – Forschungsaufenthalt
Benjamin Kaminski (University College London, GB)
March 16–21, 2021 | Research Group Meeting | https://www.dagstuhl.de/21118

21119 – Forschungsaufenthalt
Nils Jansen (Radboud University Nijmegen, NL)
March 16–21, 2021 | Research Group Meeting | https://www.dagstuhl.de/21119

21183 – Logic Programming Meets Software Language Engineering
Sebastian Erdweg (Universität Mainz, DE)
May 3–6, 2021 | Research Group Meeting | https://www.dagstuhl.de/21183

21189 – The Number of Runs in Nonbinary Strings
Johannes Fischer (TU Dortmund, DE)
May 2–7, 2021 | Research Group Meeting | https://www.dagstuhl.de/21189

21217 – Forschungsaufenthalt "Properties of Graphs Specified by a Regular Language"
Volker Diekert (Universität Stuttgart, DE)
May 25–28, 2021 | Research Group Meeting | https://www.dagstuhl.de/21217

21218 – SoftLang Koblenz Research Seminar
Ralf Lämmel (Universität Koblenz-Landau, DE)
May 25–28, 2021 | Research Group Meeting | https://www.dagstuhl.de/21218

21219 – Decomposition Methods for Two-Stage Stochastic Optimization
Bismark Singh (Universität Erlangen-Nürnberg, DE)
May 24–29, 2021 | Research Group Meeting | https://www.dagstuhl.de/21219

21224 – Semantik und KI in der Praxis
Tom Hanika (Universität Kassel, DE), Anatol Reibold (von Vicht GmbH, DE), Karl Erich Wolff (Ernst-Schröder-Zentrum Darmstadt, DE)
May 31 to June 2, 2021 | Research Group Meeting | https://www.dagstuhl.de/21224

21225 – Quantitative Software Verification
Benjamin Kaminski (University College London, GB)
May 30 to June 5, 2021 | Research Group Meeting | https://www.dagstuhl.de/21225

21233 – SoftLang Koblenz Research Seminar
Ralf Lämmel (Universität Koblenz-Landau, DE)
June 6–10, 2021 | Research Group Meeting | https://www.dagstuhl.de/21233

21255 – Workshop des GRK UnRAVeL
Joost-Pieter Katoen (RWTH Aachen, DE)
June 20–23, 2021 | Research Group Meeting | https://www.dagstuhl.de/21255

21303 – Lehrstuhltreffen Sorge
Christoph Sorge (Universität des Saarlandes, DE)
July 28–30, 2021 | Research Group Meeting | https://www.dagstuhl.de/21303

21338 – Generierung von Modellen der Laufzeit von wiss. Anwendungen auf FPGAs
Pascal Jungblut (LMU München, DE)
August 15–27, 2021 | Research Group Meeting | https://www.dagstuhl.de/21338

21339 – A Security Notion for Key-Updatable Identity-Based Signature Schemes
Sophia Grundner-Culemann (LMU München, DE)
August 15–27, 2021 | Research Group Meeting | https://www.dagstuhl.de/21339

21353 – Applications of Formal Sciences: Explainable AI
Bernhard Ganter (TU Dresden, DE), Tom Hanika (Universität Kassel, DE), Anatol Reibold (von Vicht GmbH, DE), Karl Erich Wolff (Ernst-Schröder-Zentrum Darmstadt, DE)
September 1–3, 2021 | Research Group Meeting | https://www.dagstuhl.de/21353

21369 – Forschungsaufenthalt
Anna Katharina Schmitt (Universität Mainz, DE)
September 5–10, 2021 | Research Group Meeting | https://www.dagstuhl.de/21369

21383 – DDI Cross Domain Integration for FAIR Data Sharing across Discipline Boundaries
Arofan Gregory (DDI Alliance / CODATA, US), Simon Hodson (CODATA – Paris, FR), Hilde Orten (NSD – Bergen, NO), Joachim Wackerow (GESIS – Mannheim, DE)
September 19–24, 2021 | Research Group Meeting | https://www.dagstuhl.de/21383

21393 – Interoperability for Cross-Domain Research: Use Cases for Metadata Standards
Arofan Gregory (DDI Alliance / CODATA, US), Simon Hodson (CODATA – Paris, FR), Hilde Orten (NSD – Bergen, NO), Joachim Wackerow (GESIS – Mannheim, DE)
September 26 to October 1, 2021 | Research Group Meeting | https://www.dagstuhl.de/21393

21394 – Klausurtagung der Forschungsgruppe "Dezentrale Systeme und Netzdienste"
Hannes Hartenstein (KIT – Karlsruher Institut für Technologie, DE)
September 29 to October 1, 2021 | Research Group Meeting | https://www.dagstuhl.de/21394

21403 – Klausurtagung der Graduate School "Intelligent Methods for Test and Reliability"
Ilia Polian (Universität Stuttgart, DE)
October 6–8, 2021 | Research Group Meeting | https://www.dagstuhl.de/21403

21404 – Next Steps in Body-Based Interaction
Jürgen Steimle (Universität des Saarlandes, DE)
October 4–6, 2021 | Research Group Meeting | https://www.dagstuhl.de/21404

21413 – Representation Theory and Algebraic Groups
Laura Voggesberger (TU Kaiserslautern, DE)
October 10–15, 2021 | Research Group Meeting | https://www.dagstuhl.de/21413

21414 – Klausurtagung AG Robotersysteme
Karsten Berns (TU Kaiserslautern, DE)
October 14–15, 2021 | Research Group Meeting | https://www.dagstuhl.de/21414

21423 – Retreat Software Science
Nils Jansen (Radboud University Nijmegen, NL)
October 20–22, 2021 | Research Group Meeting | https://www.dagstuhl.de/21423

21453 – Static Analyses of Program Flows: Types and Certificate for Complexity
Clément Aubert (Augusta University, US)
November 7–12, 2021 | Research Group Meeting | https://www.dagstuhl.de/21453

21469 – Forschungsaufenthalt
Friedrich Steimann (Fernuniversität in Hagen, DE)
November 15–19, 2021 | Research Group Meeting | https://www.dagstuhl.de/21469

21505 – inf-schule.de: Weiterentwicklung des elektronischen und interaktiven Lehrbuches für den Informatikunterricht
Hannes Heusel (Eduard-Spranger Gymnasium Landau, DE)
December 14–15, 2021 | Research Group Meeting | https://www.dagstuhl.de/21505

Leibniz
Gemeinschaft

Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2021 | ISSN 2199-1995
http://www.dagstuhl.de