SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2015

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2015

Herausgeber	Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH Oktavie-Allee, 66687 Wadern, Deutschland
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
Kontakt	Geschäftsstelle Schloss Dagstuhl Universität des Saarlandes Campus E1 1 66123 Saarbrücken Germany Telefon: +49 681 302–4396 Fax: +49 681 302–4397 E-Mail: service@dagstuhl.de http://www.dagstuhl.de
Redaktion	Dr. Marcel R. Ackermann, Dr. Andreas Dolzmann, Dr. Marc Herbstritt
Mitarbeit	Dr. Marcel R. Ackermann, Dr. Roswitha Bardohl, Heike Clemens, Laura Cunniff, Dr. Andreas Dolzmann, Jutka Gasiorowski, Dr. Marc Herbstritt, Oliver Hoffmann, Dr. Christian Lindig, Petra Meyer, Thomas Schillo, Christina Schwarz, Dr. Michael Wagner und die Autoren der Zusammenfassungen in Kapitel 4
Gestaltung	in Zusammenarbeit mit le-tex publishing services GmbH, Leipzig, Deutschland
Druck	Nalbacher Druckhaus, Nalbach
Druckauflage	160
ISSN	2199-1995
Online-Version	http://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	Mai 2016
Copyright	© Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH
Lizenz	Creative Commons BY-ND 3.0 DE

Leibniz-Gemeinschaft

Zertifikat seit 2012 audit berufundfamilie

Vorwort
Foreword

25 Jahre gibt es nun „Dagstuhl“. Im Juli des Berichtsjahres konnten wir dieses Jubiläum feiern. Was unter dem Namen IBFI (Internationales Begegnungszentrum für Informatik Schloß Dagstuhl) unscheinbar, aber mit großen Ambitionen 1990 begann, hat sich zu einem Informatikzentrum gebildet, um das uns international viele Wissenschaftler beneiden. Ja, das Wort „Dagstuhl“ hat sogar neue Bedeutungen bekommen. Neu gegründete, ähnliche Zentren beschreiben sich als „Dagstuhl-like“, und von Informatikerkollegen aus Übersee habe ich gelegentlich schon Aussprüche wie „Let's have a Dagstuhl“ gehört, was dann soviel bedeuten soll wie, „Lass uns einen kleinen Workshop organisieren“.

Heute firmiert das Zentrum als LZI (Schloss Dagstuhl – Leibniz-Zentrum für Informatik), steht als Mitglied der Leibniz-Gemeinschaft auf sicherer finanzieller Grundlage, hat ein wohleingespieltes Team von hervorragenden Mitarbeitern und bietet nach wie vor ein abwechslungsreiches, wissenschaftlich hoch interessantes Programm mit ein bis zwei Seminaren pro Woche über die unterschiedlichsten Themen aus der und um die Informatik. Was kann man besonders herausheben? Das Seminar im Dezember über „Graphisomorphismus“, einem seit über 30 Jahren komplexitätstheoretisch offenen Problem, bei dem der brandneue große Durchbruch von Prof. Babai, einem der Organisatoren dieses Seminars, von einem extrem fachkundigen internationalen Publikum über mehrere Tage eingehend diskutiert werden konnte. Und es gab mehrere Veranstaltungen, die mit „Artefakten“ der Informatik zu tun hatten, also Software und ganzen Umgebungen, Experimenten und Ergebnisdaten: Wie können solche Artefakte angemessen evaluiert, dokumentiert und auch archiviert werden?

Diese Artefakte beginnen auch unsere beiden anderen Abteilungen zu beschäftigen. Im Rahmen von Dagstuhl

25 years ago “Dagstuhl” got on its way. In July of this report year we celebrated this anniversary. What started in 1990 inconspicuously, but with great ambition under the name IBFI (International Meeting Center for Informatics) has become a center of computer science research that many international researchers envy us for. The word “Dagstuhl” has even assumed new meanings: newly started similar centers often describe themselves as “Dagstuhl-like”, and I have heard colleagues from oversees uttering sentences such as “Let's have a Dagstuhl”, meaning “Let us organize a small workshop”.

Today the center is known as LZI (Schloss Dagstuhl – Leibniz Center for Informatics), it is a member of the Leibniz-Association and thus in a reasonably secure funding position, it can draw on a well cooperating team of excellent staff, and of course it still offers a varied, scientifically highly interesting program of one to two seminars per week about multifarious topics in computer science and surrounding areas. What was particularly noteworthy? The seminar in December on “Graph Isomorphism”, a problem whose computational complexity had been open for over 30 years, where Prof. Babai's fresh new great breakthrough result could be discussed and checked in detail for several days by an extremely knowledgeable audience. And there were several events dealing with “artifacts” of informatics. i.e. software and entire computational environments, experiments and result data: How can such artifacts be appropriately evaluated, documented and also archived?

Such artifacts have also become a topic of interest for our other two divisions. Dagstuhl Publishing is looking at the question of a proper publication mode for such artifacts and how to provide such a mode. In the context of dblp you would like to cite such “publications” properly and to attribute them properly to authors and other contributors.

Publishing geht es darum, solchen Artefakten einen geeigneten Veröffentlichungsmodus zu finden und zu bieten. Und im Rahmen von dblp möchte man solche „Veröffentlichungen“ richtig zitieren und Autoren und andere Beitragende passend attribuieren. Wie erfolgreich diese beiden Abteilungen insgesamt sind, kann man dem Inneren dieses Jahresberichts entnehmen.

Im Jahr 2016 wird das LZI evaluiert werden. Wir sehen dieser Evaluation mit Zuversicht entgegen und freuen uns auf die kommenden Aufgaben in den nächsten Jahren.

Raimund Seidel

Im Namen der Geschäftsführung

Prof. Raimund Seidel, Ph. D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

For the overall success of these two divisions, dblp and Dagstuhl Publishing, see the inside of this report.

In 2016 LZI will be evaluated by the Leibniz-Association. We are moving towards this evaluation with confidence and look forward to our work in the coming years.

On behalf of the Managing Directors

Inhaltsverzeichnis
Contents

1 Das Zentrum Schloss Dagstuhl
Schloss Dagstuhl Center

Dagstuhls Mission

1.1

Dagstuhl's Mission

Schloss Dagstuhl – Leibniz Zentrum für Informatik erfüllt seine Mission, die Informatikforschung auf internationalem Spitzenniveau zu fördern, 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 Center for Informatic pursues its mission of furthering world class research in computer science by facilitating communication and interaction between researchers. The objective of Schloss Dagstuhl is

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

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

Entwicklung des Zentrums

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

Seit 2005 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 88 Forschungsinstituten, Bibliotheken und Museen.[1] Schloss Dagstuhl wird seit 2006 durch eine Bund-Länder-Förderung finanziert.

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

History of the Center

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

Since 2005, Schloss Dagstuhl has been a member of the Leibniz Association, a non-profit research consortium composed of 88 research institutes, libraries and museums throughout Germany.[1] Since 2006 the center is jointly funded by the German federal and state governments.

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

Seminar- und Workshop-Programm

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

Seminar and Workshop Program

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

[1] Stand Dezember 2015.
As of December, 2015.

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. Nur die besten der vom wissenschaftlichen Direktorium begutachteten Vorschläge werden in das wissenschaftliche Programm aufgenommen. Die entsprechenden Veranstaltungen finden dann durchschnittlich zwischen 6 und 18 Monaten später statt. Eine Teilnahme ist nur mit einer persönlichen Einladung durch das Zentrum möglich. Die Teilnahmegebühr ist stark subventioniert, so dass auch vielversprechende Nachwuchswissenschaftler an den Dagstuhl-Seminaren teilnehmen können. Dies eröffnet ihnen die Gelegenheit, im Zentrum mit exzellenten Experten zusammenzutreffen und neue Sichtweisen zu diskutieren.

Das Seminarzentrum ist im und rund um das 1760 erbaute Schloss Dagstuhl beheimatet und befindet sich in einer ländlichen Gegend im nördlichen Saarland, im Herzen des Dreiländerecks Deutschland, Frankreich und Luxemburg. Es bietet den Gästen eine einzigartige Arbeitsumgebung, die den Austausch mit anderen Gästen in einer wohnlichen Atmosphäre fördert. Gemütliche Sitzecken, ansprechende Essräume, eine herausragenden Informatik-Fachbibliothek, sowie eine Vielzahl von zusätzlichen Arbeits- und Freizeiträumen bieten vielfältige Möglichkeiten, damit sich die Gäste auch außerhalb des fachlichen Seminarprogramms kennenlernen und austauschen können.

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

Bibliographiedatenbank dblp

Seit 2011 betreibt Schloss Dagstuhl in enger Zusammenarbeit mit der Universität Trier die Bibliographiedatenbank dblp, welche mit mittlerweile mehr als drei Millionen Publikationseinträgen die weltweit größte, offene Sammlung bibliographischer Daten in der Informatik ist. Der dblp-Dienst ist darauf ausgerichtet, Forscher bei ihrer täglichen Arbeit zu unterstützen, etwa bei der Literaturrecherche oder beim Bezug von elektronisch verfügbaren Volltexten. Dabei gilt dblp in der Informatik insbesondere als die Referenzdatenbank für qualitätsgesicherte, normierte Bibliographiedaten. Aber auch Forschungsförderer und Entscheidungsträger unterstützt dblp, etwa durch das Pflegen und öffentlich Verfügbarmachen von personalisierten Publikationsnachweisen. Durch den Betrieb von dblp leistet Schloss Dagstuhl einen weiteren Beitrag im Rahmen seiner Mission zur Förderung der Erkennung, Verbreitung und Umsetzung neuer Informatikerkenntnisse auf international anerkanntem Niveau.

Details über dblp finden sich in Kapitel 3.

work, smaller groups of ca. 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.

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

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

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

dblp computer science bibliography

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

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

Dagstuhl Publishing

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

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 2015 — 1.2 — *News from 2015*

Das Team

Nahezu alle Mitarbeiter von Schloss Dagstuhl wurden 2015 über den Kernhaushalt des Zentrums bezahlt. Eine Ausnahme bildet das dblp-Team. Die Klaus Tschira Stiftung unterstützte Schloss Dagstuhl und dblp wie in den vergangenen Jahren auch 2015 mit einer großzügigen Spende von inzwischen 70 300 €. Zwei im Jahr 2015 neu eingestellte Mitarbeiter werden innerhalb des seit Juli 2015 laufenden SAW-Projektes „Skalierbare Autoren-Disambiguierung in Literaturdatenbanken“ aus dem Haushalt der Leibniz-Gemeinschaft finanziert.

Schloss Dagstuhl beschäftigt in der Küche zwei Auszubildende. Eine ist seit 2014 im Zentrum beschäftigt, während die andere 2015 neu eingestellt wurde. Daneben arbeiteten im Sommer 2015 zwei Praktikanten für insgesamt 6 Wochen in der IT-Abteilung. Beide arbeiteten innerhalb eines Projektes zur Erfassung der Netzwerkinfrastruktur in eine Datenbank, ihre graphischen Darstellung in einer Webanwendung sowie der Optimierung der Verkabelung. Einer der Praktikanten wurde dabei im Rahmen einer Umschulungsmaßnahme vom Arbeitsamt bezahlt.

Ende 2015 beschäftigte Schloss Dagstuhl insgesamt 34,64 Vollzeitäquivalente bzw. 49 Angestellte. Dies entspricht gegenüber Anfang des Jahres einem Zuwachs von 1,4 Vollzeitäquivalenten bzw. 4 Angestellten.

Bibliographiedatenbank dblp

Der Datenbestand der dblp computer science bibliographie ist auch im Jahr 2015 weiter stark gewachsen.

Im Laufe des Jahres wurden dem Datensatz über 367 000 neue Publikationen hinzugefügt, also mehr als 1 000 Publikationen pro Kalendertag. Dies entspricht einer erneuten Steigerung um über 4,1% gegenüber der Auf-

The Team

Nearly all staff at Schloss Dagstuhl were funded from the center's core budget in 2015. An exception is the dblp team, where – as in the previous years – one position was supported by a generous donation in the amount of 70,300 € from the Klaus Tschira Foundation. Additionally, two newly hired staff members were financed by the Leibniz-Gemeinschaft by means of the SAW project "Scalable Author Disambiguation for Bibliographic Databases".

Schloss Dagstuhl's kitchen employs two trainees. One started in 2014, whereas the other begun her training in 2015. Additionally, the IT department hosted two interns for altogether 6 weeks during the summer. They worked in a project to gather data about the network infrastructure in a database, to display it in a web front-end, and to optimize the cable system. One of the positions was financed by the employment office as part of retraining measures.

At the end of 2015, Schloss Dagstuhl had a total of 49 staff members corresponding to 34.64 full-time positions. Compared to the beginning of 2015, this is an increase of 4 staff members or 1.4 full-time equivalent positions.

dblp computer science bibliography

Between January 1st, 2015, and December 31st, 2015, the dblp database grew by more than 367,000 publication records to reach a total of more than 3.19 million records. This is another 4.1% increase when compared to the already high rate of new inclusions reached in the years before.

nahmequote des Vorjahres. Ende 2015 indexierte dblp somit bereits über 3,19 Millionen Fachartikel aus den verschiedenen Teilgebieten der Informatik.

Die Nutzung des dblp-Dienstes verstetigte sich dabei auf konstant hohem Niveau. Jeden Monat verzeichnet die dblp-Webseite bis zu 15 Millionen Seitenzugriffe von über 450 000 verschiedenen Nutzern aus aller Welt. Dies entspricht mehr als fünf Seitenzugriffen pro Sekunde, und im Durchschnitt beginnt alle drei Sekunden ein neuer Nutzer mit dblp zu arbeiten.

Das im Leibniz-Wettbewerb geförderte Projekt „Skalierbare Autoren-Disambiguierung in Literaturdatenbanken“ starte formal im Juli 2015. Das Projekt ist eine Kooperation zwischen dblp, der mathematischen Literaturdatenbank zbMATH und dem Heidelberger Institut für Theoretische Studien (HITS). Ziel ist es Verfahren zu entwickeln, um die Urheberschaft wissenschaftlicher Publikationen in Datenbanken wie dblp eindeutig zu erkennen und zuzuordnen. Das Projekt läuft bis Mitte+2018.

Mehr Informationen zu dblp finden sich in Kapitel 3.

Dagstuhl Publishing

Die Open-Access-Publikationsaktiväten haben in 2015 weiterhin starken Zuspruch bekommen, insbesondere in der Konferenzbandreihe LIPIcs, in welcher mehrere hochrangige wissenschaftliche Konferenzen, darunter z.B. ECOOP, MFCS und SNAPL.

Eine Neuerung in 2015 war die Gründung der Reihe *Dagstuhl Artifacts Series (DARTS)*, in welcher qualitätsgeprüfte Forschungsdaten und -artefakte veröffentlicht werden.

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

Öffentlichkeitsarbeit und Weiterbildung

Am 13. Mai 2015 besuchte die saarländische Ministerpräsidenten, Frau Annegret Kramp-Karrenbauer, die Stadt Wadern und kam mit einer Delegation von ca. 25 Personen nach Schloss Dagstuhl. Neben Frau Kramp-Karrenbauer gehörten der Bürgermeister, Herr Jochen Kuttler, und einige Vertreter des Kulturamtes der Stadt Wadern, zu der Delegation als auch Entscheidungsträger aus umliegenden Städten . Der wissenschaftliche Direktor von Schloss Dagstuhl, Raimund Seidel, führte die Gruppe durch die Gebäude und erklärte dabei das Dagstuhl-Konzept. Nach der geführten Tour wurde ein kleiner Empfang ausgerichtet.

Die ersten Dagstuhl-Seminare wurden im August 1990 organisiert. Seitdem ist die Anzahl der jährlich durchgeführten Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops stetig gewachsen und weitere Geschäftsfelder, wie dblp und Open Access Publishing sind hinzugekommen. Am 3. Juli 2015 feierte Schloss Dagstuhl sein 25-jähriges Bestehen mit einem Festkolloquium, bei dem der Erfolg in verschiedenen Reden hervorgehoben wurde.

Seit 1991 engagiert sich Schloss Dagstuhl im schulischen Bereich durch die Organisation einer jährlichen

Up to 15 million web pages are visited each month by more than 450,000 researchers and computer science enthusiasts all over the world. Every second, about five web pages are requested from the dblp web servers, and about every three seconds, a new user session is started.

To improve the correct attribution of publications to their unambiguous authors, a joint project of dblp, the zbMATH database for mathematical literature, and the Heidelberg Institute for Theoretical Studies (HITS) has been launched in 2015. The project had been a successful in the “Leibniz Competition 2015” and will be funded by the Leibniz Association for three years.

More information about dblp can be found in Chapter 3.

Dagstuhl Publishing

Schloss Dagstuhl’s open-access publishing services experienced an on-going strong increase in demand from the community in 2015. This was especially true of the LIPIcs conference proceedings series, which received and accepted proposals from major scientific conferences such as ECOOP, MFCS and SNAPL.

A novelty in 2015 was the establishment of the *Dagstuhl Artifacts Series (DARTS)* series in which quality-evaluated research data and artifacts are published.

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

Public Relations and Professional Training

On May 13, 2015, the prime minister of the Saarland, Mrs. Annegret Kramp-Karrenbauer, visited Wadern and Schloss Dagstuhl. A delegation of approximately 25 persons were given a brief guided tour through Schloss Dagstuhl by Dagstuhl’s Scientific Director, Raimund Seidel. The delegation included the mayor of Wadern, Mr. Jochen Kuttler, and some representatives from Wadern’s cultural office as well as some decision-makers from nearby towns. After the guided tour, a small reception was provided.

The first Dagstuhl Seminars were organized in August 1990. Since then the number of the annual Dagstuhl Seminars and Dagstuhl Perspectives Workshops has grown steadily, and other business areas, such as dblp and open access publishing were added. On July 3rd, 2015, Schloss Dagstuhl celebrated its 25th anniversary with a colloquium where the success has been highlighted in various speeches.

Since 1991, Schloss Dagstuhl has organized a training course specifically designed for computer science and mathemathics teachers working in the Saarland and the Rhineland Palatinate. The workshop is organized together with the Landesinstitut Pädagogik und Medien

Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisiert. Am 9. Dezember 2015 hat Schloss Dagstuhl die 25. Lehrerfortbildung mit einem Kolloquium am Nachmittag gefeiert.

Mehr Informationen zur Öffentlichkeitsarbeit und zu den Weiterbildungsaktivitäten finden sich in Kapitel 8.

Zusammenarbeit mit dem Heidelberg Laureate Forum

Auch im Jahr 2015 gab es wieder eine Kooperation von Schloss Dagstuhl mit dem Heidelberg Laureate Forum[2] (HLF). Diese Veranstaltung bringt herausragende Mathematiker und Informatiker, nämlich Gewinner des ACM Turing Award, des Abelpreises, der Fields-Medaille, und des Nevanlinna-Preises, mit außergewöhnlich begabten jungen Wissenschaftlern aus aller Welt zusammen. Vier ausgewählte Teilnehmer des HLF 2015 erhielten in der Woche nach der dritten Ausgabe dieses Forums die Gelegenheit zur Teilnahme an dem Dagstuhl-Seminar „Mathematical and Computational Foundations of Learning Theory“ (15361). Aufgrund des großen Erfolgs der Initiative haben alle Partner einer Fortsetzung der Zusammenarbeit für das Jahr 2016 zugestimmt.

Spender und Förderer

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

Neben zahlreichen Buchspenden durch Autoren und Gäste erhielt das Zentrum einige Exemplare des Ergänzungsbandes zum Bildband *Masters of Abstraction* von der Klaus Tschira Stiftung. Dieser Bildband enthält Schwarz-Weiß-Portraits von Preisträgern des Turingpreises, des Abelpreises, der Fields-Medaille und des Nevanlinna-Preises. Die Portraits wurden von dem Fotografen Peter Badge erstellt. Als Besonderheit enthält der Band unter anderem ein aktuelles Foto von dem seit 1991 bis zu seinem Tod 2014 völlig zurückgezogen lebenden Alexander Grothendieck.

2015 erhielt die Bibliothek von mehreren Verlagshäusern erneut zahlreiche Buchspenden. Insgesamt erhielt das Zentrum im Berichtszeitraum 795 Bände als Spende, davon mehr als 650 Monographien des Springer-Verlags. 45 Bände wurden von Gästen und Forschern der Bibliothek überlassen.

Wie in den vergangenen vier Jahren förderte die Klaus Tschira Stiftung auch in diesem Jahr die Bibliographiedatenbank dblp mit einer Spende von 70 300 €.

Das Heidelberger Institut für Theoretische Studien (HITS) hat Ende 2015 zugesagt, Dagstuhl Publishing von 2016 bis 2018 mit insgesamt 111 000 € zu unterstützen.

(LPM), Saarland, and the Pädagogisches Landesinstitut Rheinland-Pfalz (PL). On Wednesday, December 9, 2015, Schloss Dagstuhl celebrated its 25th training course for teachers with a colloquium in the afternoon.

Further details about public relations and professional training at Schloss Dagstuhl can be found in Chapter 8.

Joint Outreach with the Heidelberg Laureate Forum

2015 saw another cooperation venture between Schloss Dagstuhl and the Heidelberg Laureate Forum[2] (HLF). The HLF brings winners of the ACM Turing Award, the Abel Prize, the Fields Medal, and the Nevanlinna Prize together with exceptionally talented young scientists from all over the world. Four participants were selected and invited to participate in the Dagstuhl Seminar “Mathematical and Computational Foundations of Learning Theory” (15361), taking place during the week after the third edition of the forum. Satisfied with the outstanding success of the initiative, both partners agreed to continue the cooperation in 2016.

Sponsors and Donors

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

In 2015, the center was glad to receive numerous private book donations from guests and partners, including several copies of an addendum to the illustrated book *Masters of Abstraction* as a special gift from the Klaus Tschira Foundation. The book contains black and white photographic portraits, taken by the photographer Peter Badge, of recipients of the Turing Award, the Abel Prize, the Fields Medal, and the Nevanlinna Prize. A special feature is an up-to-date portrait of Alexander Grothendieck, who lived in isolation since 1991 until his death in 2014.

The center’s research library received a large number of book donations from several publishing houses. The number of donated volumes totaled 795, including more than 650 monographs from the Springer publishing house. 45 books were donated by guests and researchers.

As in the previous four years, Schloss Dagstuhl was grateful to receive a grant of now 70,300 € from the Klaus Tschira Foundation in support of the dblp computer science bibliography in 2015.

Just at the end of 2015, the Heidelberg Institute for Theoretical Studies (HITS) confirmed its support for Dagstuhl Publishing with 111,000 € from 2016 to 2018.

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

NSF Förderung von Nachwuchswissenschaftlern

Schloss Dagstuhl hatte im August 2012 bei der National Science Foundation (NSF) in den USA einen Antrag auf Förderung junger Wissenschaftler gestellt. Unter dem Titel „Schloss Dagstuhl – NSF Support Grant for Junior Researchers“ [3] wurde der Antrag mit Förderungsbeginn am 1. Oktober 2013 genehmigt. Für drei Jahre kann Schloss Dagstuhl bis zu 80 Nachwuchswissenschaftlern aus den USA den Besuch von Seminaren ermöglichen. Dies umfasst alle Wissenschaftler kurz vor der Promotion oder bis zu fünf Jahren nach der Promotion, sofern sie an einer US-amerikanischen Universität arbeiten. Jeder Teilnehmer erhält hierfür auf Kosten der NSF sowohl freien Aufenthalt auf Schloss Dagstuhl als auch eine Reisekostenunterstützung bis zu 750 $. Im Berichtsjahr konnte dadurch 19 Forschern aus den USA eine Teilnahme an insgesamt 17 Seminaren ermöglicht werden.

NSF Grant for Junior Researchers

In August 2012, Schloss Dagstuhl applied to the National Science Foundation (NSF), USA, for support for junior researchers. The application was approved, and the NSF Grant for Junior Researchers was established, effective from October 1, 2013. Over a period of three years, up to 80 junior researchers from the US will thus be able to attend Dagstuhl Seminars. Included are researchers working at US-American universities who are soon to complete their doctorate or have completed it within the last five years. Each participant receives a free stay as well as a travel stipend in the amount of $ 750 at the expense of the NSF. In 2015, 19 US-based scientists were supported and hence able to participate in overall 17 Dagstuhl Seminars.

Vereinbarkeit von Beruf und Familie

Schloss Dagstuhl hatte bereits 2012 das Audit der berufundfamilie Service GmbH erfolgreich durchlaufen und die drei Jahre gültige Zertifizierung „audit berufundfamilie“ erhalten. In dem Auditierungsprozess wird sichergestellt, dass Firmen eine familien- und lebensphasenbewusste Personalpolitik beibehalten oder in einem systematischen Prozess entwickeln. Mit der Bestätigung des wiederum drei Jahre gültigen Zertifikats am 31. Mai 2015 hat Schloss Dagstuhl die erste Re-Auditierung erfolgreich abgeschlossen.

Compatibility of Family and Career

As early as 2012, Schloss Dagstuhl was successfully audited by berufundfamilie Service GmbH and received the certificate “audit berufundfamilie” (family and career) which is valid for three years. The auditing process is intended to ensure that the companies’ human resource policy continues to be family-conscious and adapted to the different phases of life, or that such a policy is developed in a systematic process. The certificate was confirmed and extended for another three years on May 31, 2015, meaning that Schloss Dagstuhl successfully passed through the first re-audit.

Baumaßnahmen und Renovierung

Schloss Dagstuhl hält die verschiedenen Gebäuden des Zentrums laufend instand und mordernisiert diese. So wurden im Sommer 2015 als größtes Projekt ein Teil der Fensterfronten im 1993 fertig gestellten Neubau saniert. Große Flächen, wie z. B die Verglasung des sogenannten Kreuzgangs zum Innenhof, die Fenster im Eingangsbereich und im Bereich des Bibliotheksturms wurden dabei vollständig erneut. Im Gegensatz zu den früheren Holzfensterrahmen kamen nun Aluminiumrahmen zum Einsatz, die eine deutlich höhere Lebenserwartung haben. Die bisherige Verglasung wurde durch eine spezielle Wärmeschutzverglasung mit integriertem Sonnenschutz ersetzt.

Nachdemin den letzten Jahren bis einschließlich 2015 die Bäder im Neubau teilmodernisiert wurden, wurde 2015 auch im Anbau des Schlosses ein Badezimmer komplett renoviert und neu gestaltet. An diesem Muster werden nun Erfahrungen gesammelt, wie sich die neue Einrichtung sowohl für Gäste als auch bei der täglichen Reinigung durch unseren Hausdienst bewährt.

Im neuen Gästehaus wurden vor die große Fensterfront der Gästezeimmer Jalousien montiert. Diese dienen der Beschattung und damit auch dem Schutz vor dem Aufhei-

Construction Work and Renovation

Schloss Dagstuhl continually maintains and modernises all of the centre‘s buildings. The biggest project in 2015 was the partial renovation of the windows in the New Building from 1993. Large surfaces such as the glazing of the so-called cloister by the courtyard as well as the windows in the entrance area and in the vicinity of the library tower were replaced completely. The new window frames are made of aluminum instead of wood, thus having a significantly longer lifespan. The former glazing was replaced by a special thermal protection glazing with integrated solar protection. After the New Building’s bathrooms had been partially renovated in the previous years including 2015, one bathroom in the main building was 2015 completely renovated and redesigned. It will serve as a prototype in order to see whether the alterations prove successful with guests and cleaning staff. 2015 also saw the installation of sun-blinds on the guesthouse’s window façade. They shade the rooms, preventing them from overheating due to excessive sun exposure. The step proved successful: long-term measurements showed that the room temperature decreased by as much as 5 degrees after the blinds were installed.

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

zen der Räume durch übermäßige Sonneneinstrahlung. Die Maßnahme war sehr erfolgreich: Messungen vor und nach dem Einbau haben ergeben, dass bei Sonneneinstrahlung mit heruntergelassenen Jalousien die Temperatur in den Zimmern um bis zu 5 Grad niedriger blieb als vorher.

Ausstattung

Neben diesen größeren Maßnahmen und Neuerungen hat Schloss Dagstuhl auch darauf geachtet, weitere Maßnahmen zur Verbesserung des Komforts und Ambiente umzusetzen. Nachdem in den letzten Jahren der Speisesaal bereits neu bestuhlt wurde, wurden dort 2015 auch alle Tische im gewohnten Stil des Raums ersetzt. Durch eine spezielle Nano-Beschichtung der Oberflächen wurde eine erhöhte Strapazierfähigkeit erreicht. Weiterhin wurden die Stühle auf den Zimmern des Gästehauses, des Neubaus und in der Bibliothek durch neue ersetzt.

Um unsere Open-House Politik nicht aufgeben zu müssen, aber dennoch dem Sicherheitsbedürfnis unserer Gäste Rechnung zu tragen, wurden Anfang 2015 in allen Zimmern Safes eingebaut. Diese können mit einer selbst gewählten Nummer verschlossen werden und sind groß genug, um auch einen Laptop darin aufzubewahren.

Im Hörsaal „Saarbrücken“ wurde vor der mittleren Tafel eine Leinwand montiert, so dass nun eine größere Projektionsfläche insbesondere für Full-HDMI zur Verfügung steht. Trotz Leinwand können die Tafeln links und rechts der Leinwand voll genutzt werden.

IT Services

Die IT hat insbesondere in der Netzwerktechnik weitere Maßnahmen ergriffen, um den Zugriff der Gäste und Mitarbeiter auf das Internet sicherzustellen. So wurde die Anbindung an die beiden von der DFN betriebenen Leitungen optimiert. Ein Firewall/Router-Cluster mit Failover- und Lastverteilungsfunktion wurde neu beschafft und seitens der IT den spezifischen Bedürfnissen von Schloss Dagstuhl gemäß programmiert.

Im Rahmen der Fenstersanierung wurden auch alle WLAN-Zugriffspunkte im Neubau neu verteilt, weitere montiert und alle Geräte neu verkabelt, um die Empfangsmöglichkeit in den Gästeräumen weiter zu optimieren.

Die Hausautomatisierung auf Basis der HomeMatic-Installation wurde weiter ausgebaut. 2015 wurden vor allem Elemente zur automatischen Lichtsteuerung in Betrieb genommen.

Die Web-Oberfläche zum Einreichen der Beiträge für die Dagstuhl Reports wurde 2015 komplett von der IT-Abteilung reimplementiert. Das neue System bietet ein weitergehendes Rechtesystem, Anschluss an die zentrale Verwaltung der Gastdaten und eine moderne und komfortable Oberfläche.

Die IT Abteilung hat aufbauend auf dem Raspberry Pi Kleinstcomputer eine Reihe von Projekten verwirklicht, z.B. wurde eine Überwachung der Stromversorgung, die Ausfälle automatisch per Mobilfunk meldet, realisiert. Ein weiteres Projekt war die Steuerung der Projektion der Informationsmaterialien auf dem großen Touchscreen

Facilities

Aside from these large-scale measures and alterations, Schloss Dagstuhl took further steps to modernise its facilities for the purpose of comfort and atmosphere. After the dining room chairs had already been replaced in recent years, now all tables were replaced as well, re-establishing the usual decor of the room. The durability of the table surfaces is increased by a specific nano-coating. Plus, the chairs in the rooms of the guesthouse and New Building as well as in the library were replaced. In early 2015, all rooms were equipped with safes, thus making it possible for us to maintain our open house policy while catering to our guests' safety need at the same time. The safes can be locked with a self-chosen number and are big enough for storing laptops. The seminar room "Saarbrücken" was equipped with a screen in front of the middle blackboard, providing a bigger projection surface, especially for Full HDMI. Despite the screen, the blackboards left and right of it can be used as before.

IT Services

Especially with regard to network technology, the IT department took further steps to guarantee internet access for guests and staff. For instance, the connection to the two DFN-operated lines was optimized. A firewall/router cluster that provides for failover and load distribution was acquired and programmed according to Schloss Dagstuhl's specific requirements.

As part of the window renovation, all WLAN access points in the New Building were redistributed, new ones were installed and all devices were rewired, all of which improves the internet access in the guest rooms.

Schloss Dagstuhl's HomeMatic house automation was expanded. 2015 mainly saw automatic lighting control elements being put into operation.

The web interface used for submitting contributions to the Dagstuhl Reports was entirely re-implemented by the IT department in 2015. The new system includes an extended rights management, access to the main guest data management and a modern, user-friendly interface.

Building on the Rasperry Pi microcomputer, the IT team realised several projects. For instance, the power supply is now monitored and failures are reported automatically via mobile radio. Another project was the projection control for information on the big touchscreen next to the seminar room "Saarbrücken" based on the Raspberry Pi. Amongst other things, the IT department uses this computer with a color LED dot-matrix display for signaling incoming tickets as part of the IT helpdesk system.

neben dem Hörsaal „Saarbrücken“ aufbauend auf einem Raspberry Pi. Intern benutzt die IT einen solchen Rechner mit einer farbigen LED-Matrixanzeige für die Signalisierung eintreffender Tickets als Teil des Helpdesk-Systems.

Bibliothek

Seit Juli 2015 bietet Schloss Dagstuhl seinen Gästen uneingeschränkten Zugriff auf die digitale Bibliothek „IEEE Xplore“, d.h. auf alle Volltexte, die vom *Institute of Electrical and Electronics Engineers* (IEEE) elektronisch angeboten werden. Unter der Verhandlungsführung der TIB Hannover konnte für interessierte Institute der Leibniz-Gemeinschaft dieser neue Lizenzvertrag abgeschlossen werden.

Als Ergebnis des Arbeitskreises „Bibliotheken der Leibniz-Gemeinschaft“ hat die Leibniz-Gemeinschaft über die TIB Hannover mit der Firma EBSCO Information Services einen Rahmenvertrag für den Zugriff auf die Metasuchmaschine „EBSCO Discovery Service“ (EDS) abgeschlossen. Leibniz-Institute können über diesen EDS zu den im Rahmenvertrag vereinbarten Bedingungen lizenzieren. Schloss Dagstuhl hat EDS zum Januar 2016 lizenziert.

Die noch bis 2015 in der Bibliothek verwendeten „Regeln für die Alphabetische Katalogisierung“ (RAK) wurden durch ein neues Regelwerk abgelöst: Den „Resource Description and Access“ (RDA) Standard. Seit Januar 2016 ist RDA die verbindliche Grundlage für die Katalogisierung im Südwestdeutschen Bibliotheksverbundes (SWB).

Library

Since July 2015, all guest of Schloss Dagstuhl enjoy unrestricted access to the “IEEE Xplore” and “IEEE/IET Electronic Library” (IEL) digital libraries, that is, full access to all scholarly articles published by the *Institute of Electrical and Electronics Engineers* (IEEE). This has been made possible by a new licensing agreement between IEEE and the Leibniz Association (represented by TIB Hanover).

As a result of the Libraries Working Group of the Leibniz Association, a new licensing agreement could be reached between the Leibniz Association (represented by TIB Hanover) and EBSCO Information Services. Schloss Dagstuhl licensed access to the “EBSCO Discovery Service” (EDS) meta search engine, starting in January 2016.

The RAK catalog standard, which had been used by the Dagstuhl Library until the end of 2015, has been retired and replaced by the new “Resource Description and Access” (RDA) catalog standard. Starting 2016, the RDA standard has become mandatory in the SWB network of libraries in Baden-Württemberg, Saxony, Saarland, and Rhineland-Palatinate.

2 Das wissenschaftliche Seminar-Programm 2015
Scientific Seminar Program 2015

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 22 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.

Für jedes Dagstuhl-Seminar wird ein Dagstuhl Report erstellt, der eine Zusammenfassung des Seminarverlaufs, eine Kurzübersicht über die gehaltenen Vorträge und eine Zusammenfassung grundsätzlicher Ergebnisse enthält. Der Bericht gewährleist eine zeitnahe Kommunikation der Ergebnisse. Die Zeitschrift *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.[4]

In den Kapiteln 7 und 15 sind alle Veranstaltungen, die 2015 stattfanden, aufgelistet, zusammen mit Zusammenfassungen der Seminare und Perspektiven-Workshops. Auf der Dagstuhl-Website ist das Seminar-Programm für die kommenden 24 Monate verfügbar.

2.1 *Dagstuhl Seminars*

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

Participation in these events – which generally last one week – is possible only by way of personal invitation from Schloss Dagstuhl. The center assumes part of the associated costs in order to enable the world's most qualified scientists, including young researchers and doctoral students, to participate. Among Dagstuhl's guests have been 22 winners of the ACM Turing Award, the highest achievable award within the international computer science community.

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

Each Dagstuhl Seminar is asked to contribute a record of the seminar proceedings in the form of a Dagstuhl Report. The report gives an overview of the seminar's program, talks, and results in a journal-like manner to allow for a high visibility and timely communication of its outcome. The periodical *Dagstuhl Reports* is published in one volume with 12 issues per year; each issue documents the Dagstuhl Seminars and Dagstuhl Perspectives Workshops of a given month. Dagstuhl Reports are openly accessible and can be downloaded at the Dagstuhl website.[4]

Chapters 7 and 15 provide a comprehensive list of all events that took place during the year under review and summaries of the 2015 Seminars and Perspectives Workshops. A seminar program covering the coming 24 months is available on the Dagstuhl website.

[4] http://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.

Die Dagstuhl-Perspektiven-Workshops, die 2015 statt fanden, sind in Fig. 2.1 aufgelistet.

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

Eine Liste der vergangenen und kommenden Dagstuhl-Perspektiven-Workshops ist auf der Dagstuhl-Website verfügbar.[6]

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
- trigger innovation processes

Dagstuhl Perspectives Workshops held in 2015 are listed in Fig. 2.1.

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

Past and upcoming Dagstuhl Perspectives Workshop can be found on our web site.[6]

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

Digital Scholarship and Open Science in Psychology and the Behavioral Sciences http://www.dagstuhl.de/15302
Power-Bounded HPC Performance Optimization http://www.dagstuhl.de/15342
Present and Future of Formal Argumentation http://www.dagstuhl.de/15362
Artifact Evaluation for Publications http://www.dagstuhl.de/15452

Fig. 2.1
Dagstuhl Perspectives Workshops held in 2015.

2.3 Einreichung der Anträge und Begutachtungsverfahren

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

Das Zentrum erbittet zweimal im Jahr Themenvorschläge von führenden Wissenschaftlerinnen und Wissenschaftlern aus der ganzen Welt, die ihre Seminaranträge zusammen mit einer vorläufigen Teilnehmerliste einreichen. Die Anträge werden dann vom Wissenschaftlichen Direktorium begutachtet und abschließend bei zweitägigen Sitzungen auf Schloss Dagstuhl intensiv diskutiert und 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 die Miteinbeziehung junger und weiblicher Wissenschaftler geachtet.

Die Informatikforscher zeigten 2015 wieder ein hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 99 Anträgen in den Antragsrunden im Januar und Juni 2015. Etwa 66 % der eingereichten Anträge wurden genehmigt. Der Großteil der Anträge genügten den Antragskriterien überdurchschnittlich. Ablehnungen basierten vor allem auf dem Ziel der Erstellung eines thematisch ausgewogenen Seminar-Programms, sowie der starken Konkurenz in Feldern mit einer Häufung von Anträgen. In den vergangenen 7 Jahren variierte die Rate der angenommen Anträge zwischen 66 % und 76 % (siehe Fig. 2.2).

Unter den 65 in 2015 neu genehmigten Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops gab es wie in den vergangenen Jahren wieder verschiedene Konstellationen bzgl. Dauer und Größe (vgl. Fig. 2.3). Von diesen konnten 15 Seminare bereits 2015 ausgerichtet werden, der Großteil wurde jedoch für das Seminar-Programm in 2016 eingeplant.

2.3 *Proposal Submission and Review Process*

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

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

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

The international scientific community expressed a lively interest in organizing seminars and workshops at Schloss Dagstuhl in 2015, submitting 99 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2015 and June 2015 submission rounds. The quality of the proposals was excellent, resulting in a 66 % acceptance rate by Dagstuhl's Scientific Directorate. Since 2008, proposal acceptance rates have tended to range between 66 % and 76 % (see Fig. 2.2).

Among the 65 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2015 there is – as in the past years – a wide variation with regard to length and size (see Fig. 2.3). Most of these seminars are part of the 2016 seminar program, although it was possible to schedule 15 of them already in 2015.

2.4 Seminar-Programm 2015

Seit 2012 ist es aufgrund des neuen Gästehauses möglich, zwei Seminare parallel in einer Woche zu veranstalten. Typischerweise werden ein großes und ein kleines Seminar

2.4 *The Seminar Program in 2015*

Since the new guest house opened in 2012, it has been possible for the center to schedule two parallel seminars – typically a large one and a small one – in any given week.

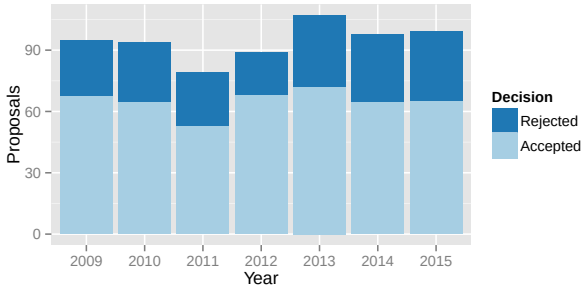

Fig. 2.2
Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2009–2015.

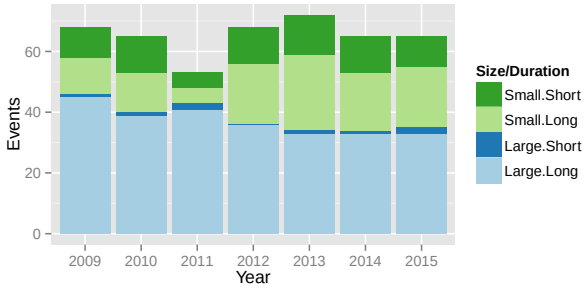

Fig. 2.3
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2009–2015. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

zusammengelegt. In 2015 fanden in 30 von 48 Wochen Seminare parallel statt, was 62 % der verfügbaren Wochen entspricht. Im Vergleich mit den Zahlen aus 2014 (69 %), 2013 (62 %), 2012 (35 %) und 2011 (12 %) lässt sich feststellen, dass eine gewisse Sättigung eingetreten ist, die sich nahe an der optimalen Auslastung befindet. In Fig. 2.4 ist die Entwicklung der vergangenen Jahre dargestellt.

Dadurch, dass nun zwei Seminare parallel veranstaltet werden können, ist in den letzten Jahren die Gesamtanzahl an Seminaren pro Jahr gestiegen. Bedingt durch die Sanierung der Fenster im Neubau fanden mit 72 Seminaren in 2015 ein paar Seminare weniger statt als im Vorjahr.

In 2015, there were two parallel seminars per week on 30 out of 48 weeks, or on 62 % of all available weeks. This rate indicates a near-optimal saturation state, as compared against the figures 2014 (69 %), 2013 (62,%), 2012 (35 %), and 2011 (12 %). Fig. 2.4 shows the evolution in recent years.

The scheduling of parallel seminars has had the effect of increasing the overall number of seminars at Schloss Dagstuhl in recent years. This number reached 72 in 2015 a little bit less than in the previous year caused mainly by the window reconstruction in summer.

Angaben zu Teilnehmern und Organisatoren

2.5 *Participant and Organizer Data*

Viele der internationalen Teilnehmer der Seminare waren schon öfter in Dagstuhl. Dennoch zieht das Zentrum jedes Jahr auch neue Gesichter an, was den ständigen Wandel in der Forschung widerspiegelt. So nahmen – wie in den Vorjahren auch – in 2015 knapp die Hälfte (46 %, 1 075 von 2 316) der Wissenschaftler das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während weitere 20 % der Wissenschaftler an nur einem Seminar in den Jahren vorher teilgenommen hatten. Ein wenig andere Zahlen leiten sich aus unserer Gastumfrage ab. Hier ergibt sich, dass etwa 40 % der Antwortenden 2015 das erste Mal und weitere 17 % zum zweiten Mal (siehe Fig. 2.5a) teilgenommen haben.

Participants in Dagstuhl Seminars come from all over the world and a significant number of them choose to repeat the experience. Nevertheless, we see many fresh new faces every year, reflecting the changing informatics research across the globe. As in the previous year, also in 2015 a bit less than the half (1,075 of 2,316, or 46 %) researchers were first-time visitors to Dagstuhl. About 20 % additional researchers had already attended one previous seminar in the years before. Slightly different numbers are obtained from our guest survey: About 40 % of the reponders were first-time visitors and additional 17 % states their second visit (see Figure 2.5a).

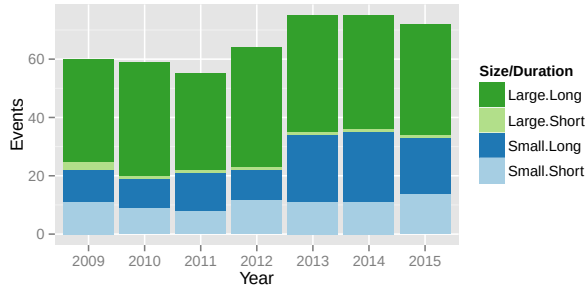

Fig. 2.4
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2009–2015. Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

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 Seminare und Workshops in 2015, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein (siehe Fig. 2.5b). Diese ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern ist im Laufe der Jahre relativ konstant geblieben, was die Bemühungen des Zentrums zur Aufrechterhaltung der „Dagstuhl-Verbindung“ zwischen herausragenden jungen Wissenschaftlern und ihren erfahrenen Kollegen zeigt.

Mit 73 % war der Anteil von Gästen aus dem Ausland 2015 erneut sehr hoch. Das Diagramm in Fig. 2.5c zeigt die regionale Verteilung der Gäste für 2015 bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops. Mehr Details kann Kapitel 14 entnommen werden.

In 2015 war mehr als die Hälfte aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt (siehe Fig. 2.6a). Der Anteil an weiblichen Seminarteilnehmern war mit 15 % wieder erfreulich hoch (siehe Fig. 2.6b).

A healthy number of these guests were young researchers at the start of their careers, for whom the Dagstuhl experience can be of lifelong value. Approximately 35 % of 2015 seminar and workshop survey respondents self-classified as junior (see Fig. 2.5b). This proportion of junior to senior researchers has remained relatively constant over the years, reflecting the center's determined effort to maintain the "Dagstuhl connection" between brilliant junior scientists and their senior colleagues.

At over 73 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2015. The chart in Fig. 2.5c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2015. For a detailed breakdown please refer to Chapter 14.

In 2015, more than the half of all organizer teams in our scientific seminar program were mixed with respect to gender (see Fig. 2.6a). The percentage of female seminar participants was also high both in total and relative terms, at 15 % (see Fig. 2.6b).

Themen und Forschungsgebiete 2.6 *Topics and Research Areas*

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

Im Folgenden sind beispielhaft einige thematische Schwerpunkte und dazugehörige Seminare aufgeführt. Die Aufzählung der Themen und Seminare hat keinen Anspruch auf Vollständigkeit und ist lediglich ein Versuch, einen kurzen Einblick in das umfangreiche Seminar-Programm zu geben. Kapitel 7 bietet mit den Kurzzusammenfassungen der Seminar- und Workshops einen vollständigen

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

The following overview gives some topical focal points and seminars from 2015. 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 2015. The seminar summaries in Chapter 7 provide a full overview of the 2015 scientific seminar program.

Among the seminars which addressed topics from theoretical computer science, there was a special one that

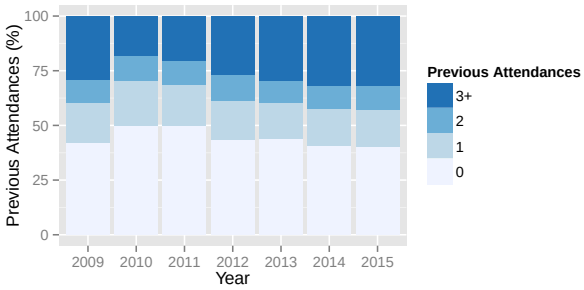

(a) Distribution of the number of previous attendances of participants, according to survey data.

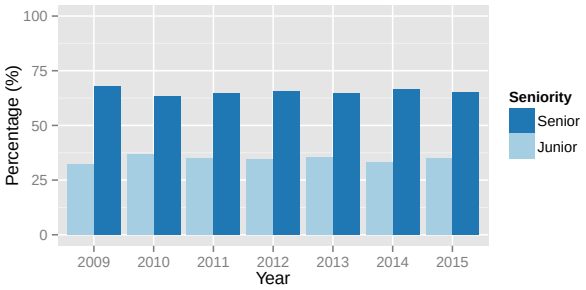

(b) Percentage of junior researchers, according to survey data.

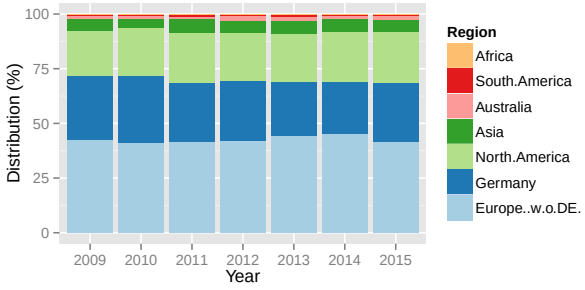

(c) Distribution of the origin of participants by region.

Fig. 2.5
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2009–2015.

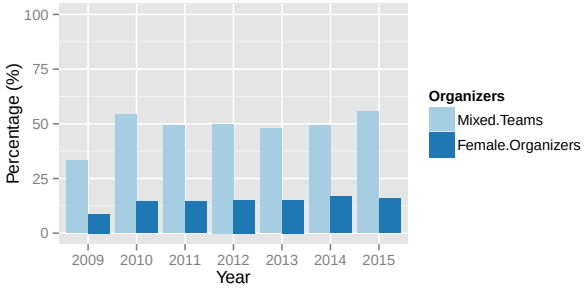

(a) Female organizers and mixed-gender organizer teams.

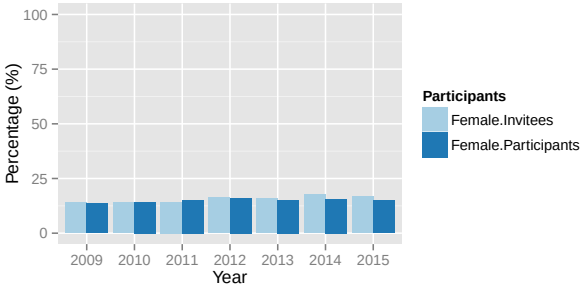

(b) Female invitees and participants.

Fig. 2.6
Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2009–2015.

Überblick über das wissenschaftliche Seminar-Programm des Jahres 2015.

Unter den Seminaren, die sich Themen aus der theoretischen Informatik gewidmet haben, gab es zum Jahresende mit dem Dagstuhl-Seminar 15501 *The Graph Isomorphism Problem* etwas ganz Besonderes: Laszlo Babai, einer der Co-Organisatoren des Seminars, hatte kurz vor Seminarbeginn bekannt gegeben, einen Beweis gefunden zu haben, dass das Graphenisomorphieproblem in quasipolynomieller Zeit gelöst werden kann. Zum Hintergrund: Aus komplexitätstheoretischer Sicht gehört das Graphenisomorphieproblem zu den wenigen Problemen, die zwar in der Klasse NP liegen, aber es unklar ist (unter der Annahme das P≠NP), ob es zur Klasse P oder etwa zur Klasse der NP-vollständigen Problemen gehört. Mit der Existenz eines quasipolynomiellen Algorithmus wurde nun die Tür aufgestoßen in Richtung Zugehörigkeit zur

took place at the end of the year: It was Dagstuhl Seminar 15501 *The Graph Isomorphism Problem*. Laszlo Babei, one of the co-organizers of the seminar, announced shortly before the seminar that he has found a proof that the Graph Isomorphism problem can be solved in quasipolynomial time. For explanation: The Graph Isomorphism problem is a problem that belongs to the complexity class NP. It is among the few problems for which it is neither known whether it also belongs to class P, nor whether it is NP-complete (assuming P≠NP). Now, with the existence of a quasipolynomial algorithm for solving the Graph Isomorphism problem, Laszlo Babai opened the door towards the problem being in the class P. Hence, theoretical computer scientists worldwide judged this as a groundbreaking result. But interestingly, Laszlo Babai also announced that he is going to discuss his proof in detail with experts from his community, before submitting the proof to a scholarly

Klasse P. Daher wurde das Ergebnis von Laszlo Babai in der Fachwelt umgehend als bahnbrechend bewertet. Laszlo Babei hatte aber ebenfalls angekündigt, die Korrektheit seines Beweises durch ausgiebige Diskussionen mit den Fachkollegen zu untermauern, bevor er den Beweis zur Begutachtung bei einer Fachzeitschrift einreicht. Das Dagstuhl-Seminar 15501 bot hierfür die perfekte Grundlage. Anuj Dawar, ebenfalls Co-Organisator des Seminars, hat in einem nachfolgenden Bericht[7] die Diskussionen in Dagstuhl wie folgt beschrieben: "*At the seminar, we scheduled four hours of talks, spread over the first two afternoons for Laci Babai to present the main ideas involved in his proof. It was an intense blackboard presentation given with great enthusiasm. In the event, in response to demand from participants for more details, Laci gave an extra hour-long unscheduled presentation on Wednesday afternoon, after the traditional Dagstuhl excursion. The paper containing the full details of the proof*[8] *was released on the arXiv on the first day of the seminar, giving participants the opportunity to consult it for details and it generated vivid discussion.*" Es war mucksmäuschenstill während Laszlo Babei seine Vorträge gehalten hat und die Spannung, die in der Luft lag, war überall in Dagstuhl zu spüren an diesen Tagen.

Die fortschreitende digitale Vernetzung der Welt gab Anlass für mehrere Seminare im Bereich *Netzwerke*. Zum einem mit klassischen Themen in Bezug auf das Routing von Datenpaketen (15102 – *Secure Routing for Future Communication Networks*) oder verteiltes Rechnen (15072 – *Distributed Cloud Computing*), aber insbesondere auch mit Themen, die die Notwendigkeit, Netzwerke theoretisch untersuchen zu können, um allgemeingültige Aussagen formal beweisen zu können sowie Eigenschaften formal sicherstellen zu können, zum Inhalt hatten (15071 – *Formal Foundations for Networking*, 15112 – *Network Calculus*). Die Annäherung von Praxis und Theorie in diesem Bereich war längst überfällig und hat durch die Dagstuhl-Seminare deutlichen Anschub erhalten.

Diese kleine Auswahl von Seminaren soll aber nicht darüber hinwegtäuschen, dass jedes der in 2015 veranstalteten Seminare wichtige Themen adressiert hat, die von den beteiligten Wissenschaftler mit großem Engagement diskutiert wurden und so die weitere Entwicklung in den einzelnen Gebieten wieder ein gutes Stück weitergebracht hat.

journal, simply to ensure that the proof is correct. Dagstuhl Seminar 15501 was the perfect spot for this assessment. Anuj Dawar, also a co-organizer of the seminar, wrote a report[7] subsequently in which he described the discussion at Dagstuhl as follows: *At the seminar, we scheduled four hours of talks, spread over the first two afternoons for Laci Babai to present the main ideas involved in his proof. It was an intense blackboard presentation given with great enthusiasm. In the event, in response to demand from participants for more details, Laci gave an extra hour-long unscheduled presentation on Wednesday afternoon, after the traditional Dagstuhl excursion. The paper containing the full details of the proof*[8] *was released on the arXiv on the first day of the seminar, giving participants the opportunity to consult it for details and it generated vivid discussion.*" It was as quiet as a mouse during the lectures of Laszlo Babei and the atmosphere at Dagstuhl was full of suspense during these days.

The ever-increasing digital interconnectedness gave the motivation for several seminars about *Networking* topics. One the one hand with seminars addressing classical issues like packet routing (15102 – *Secure Routing for Future Communication Networks*) and distributed computing (15072 – *Distributed Cloud Computing*). On the other hand with seminars that looked at the formal foundations which are required to analyze networks formally such that general propositions can be proved and such that properties can be formally ensured (15071 – *Formal Foundations for Networking*, 15112 – *Network Calculus*). The convergence of practice and theory in this area was overdue and the seminars gave substantial impetus.

This brief selection of seminars should not hide the fact that each of the 2015 seminars addressed important topics which were discussed by the involved researchers with great commitment and hence pushed forward the development in the individual areas.

2.7 Weitere Veranstaltungstypen

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 und in Kooperation mit der GI durchgeführt und von der GI sowie von Dagstuhl gefördert werden

2.7 *Further Event Types*

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

- GI-Dagstuhl seminars, sponsored by the German Informatics Society (GI) in association with Schloss Dagstuhl, that bring young scholars together to discuss and learn about a specific topic

[7] http://icetcs.ru.is/luca/anuj-report.pdf
[8] http://arxiv.org/abs/1512.03547

- Sommerschulen, Weiterbildungsveranstaltungen, Lehrerfortbildungen, Ausbildung von jungen Journalisten und Volontären
- Klausurtagungen von Graduiertenkollegs, GI-Fachgruppen und anderen akademischen und industriellen Arbeitsgruppen
- Forschungsaufenthalte

Das Angebot, Dagstuhl zu einem wissenschaftlichen Forschungsaufenthalt zu besuchen, wird regelmäßig genutzt. In den meisten Fällen sind es Einzelpersonen, die sich für eine oder mehrere Wochen für intensive Studien nach Dagstuhl in Klausur zurückziehen.

- summer schools, continuing education courses sponsored by the German Informatics Society (GI), vocational training for teachers and instructors, and educational and training workshops for young journalists and trainees
- conferences of graduate research training groups, GI specialist groups and other academic and industrial working groups
- research stays

People regularly take advantage of Dagstuhl's offer to use the center for research stays. In most cases these are individuals 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 inhaltlichen und organisatorischen Aspekten nach der Zufriedenheit ihres Besuchs. Die Ergebnisse jedes Fragebogens werden im Haus wöchentlich allen Abteilungen zugänglich gemacht, um eine schnelle Reaktion auf Probleme und Wünsche zu erreichen. Gleichzeitig werden die anonymisierten Ergebnisse von inhaltlichen Fragen den Teilnehmern eines Seminars per E-Mail mitgeteilt, typischerweise in der Woche nach ihrem Aufenthalt. So erhalten insbesondere Organisatoren Rückmeldungen über den Verlauf des Seminars und Hinweise für die Organisation von zukünftigen Seminaren. Seit 2013 werden diese statistischen Ergebnisse mit Hilfe von aussagekräftigen Diagrammen aufbereitet und als PDF-Dokumente zur Verfügung gestellt.

Fig. 2.7 zeigt die Zufriedenheit dieser Teilnehmer im Jahr 2015 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 467 Fragebögen, welche die Meinung von etwa 59 % der 2 474 Teilnehmer repräsentieren. Das durchweg sehr gute Ergebnis ist Anerkennung und Herausforderung zugleich.

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

Quality Assurance

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

Fig. 2.7 shows the satisfaction of responding participants in 2014 with regard to selected aspects of their stay. The results were compiled from 1,467 questionnaires, representing the responses of about 59 % of all participants (2,590). These excellent results are not only a recognition of the center's past work but also pose a challenge to its future work.

Since 2013, Schloss Dagstuhl has also been offering all Dagstuhl Seminar organizers a more transparent invitation process by giving them direct access to the status of invitee replies via a dedicated webpage. The page is available 24/7 and has met with very positive feedback from organizers.

Auslastung des Zentrums 2.9

Auch 2015 konnte Schloss Dagstuhl die durch das neue Gästehauses ermöglichte hohe Auslastung weitgehend halten. Insgesamt gab es 2015 mit 16 198 Gasttage, wobei 13 344 Gasttage auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Bezogen auf die Seminar- und Workshopgäste bedeutet dies einen geringen Rückgang verglichen mit 2015. Es fanden im Berichtsjahr

Utilization of the Center

Thanks to the new guest house, Schloss Dagstuhl was able to uphold the high capacity utilization again in 2015. There were 16,198 overnight stays in total, with 13,344 overnight stays in seminars and perspective workshops. The latter was slightly less than in 2014. The center hosted a total of 107 events with 3,298 guests in 2015. See Chapter 14 for further details.

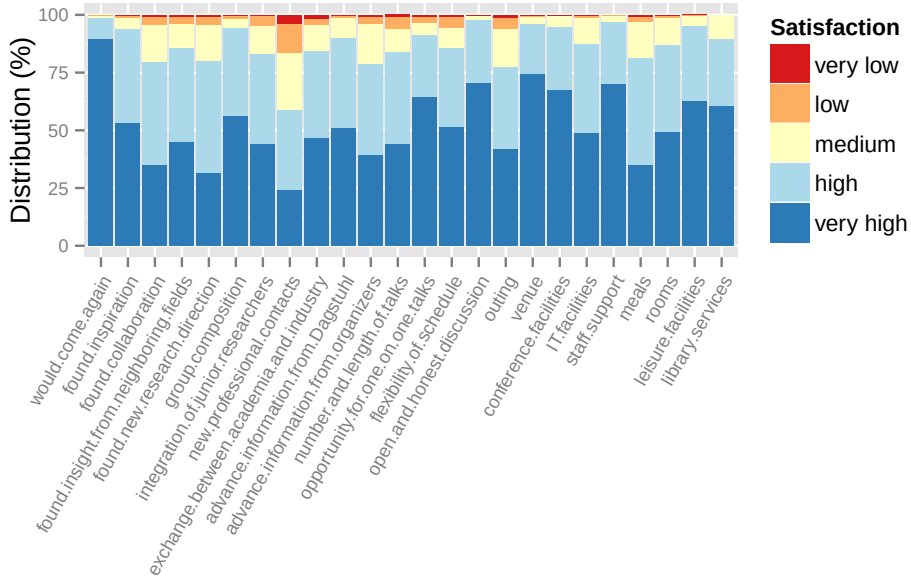

Fig. 2.7
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2015. According to survey results.

107 Veranstaltungen mit insgesamt 3 298 Gästen statt. Weitere Details können Kapitel 14 entnommen werden.

Die Wochenenden blieben 2015 ebenso unbelegt wie jeweils zwei Wochen im Juli/August und am Jahresende. Diese wurden zu Instandhaltungs- und Verwaltungsarbeiten benötigt.

Ein umfassendes Verzeichnis aller Veranstaltungen auf Schloss Dagstuhl im Jahr 2015 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 15. Auf unserer Webseite ist ein Kalender[9] verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

Weekends were kept free in 2015, as well as two weeks in July/August and at the end of the year, this time being required for maintenance work to building facilities and administrative work.

A comprehensive listing of all events at Schloss Dagstuhl in 2015, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and host-only events such as meetings and summer schools can be found in Chapter 15. See the Schloss Dagstuhl website to view our calendar[9] of upcoming events and further information and material on all events past, present and future, e.g. aims and scope, participant list, and concluding report.

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

3 Bibliographiedatenbank dblp
dblp computer science bibliography

3.1 Offene Bibliographiedaten für die Informatik

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

Die „dblp computer science bibliography“ leistet auf diesem Gebiet nun bereits seit über 20 Jahren einen substanziellen Beitrag durch die offene Bereitstellung qualitätsgeprüfter und aufbereiteter Publikationsdaten für die gesamte Informatik. Dabei unterstützt dblp die Informatik-Forschung auf gleich mehreren Ebenen, etwa durch:

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

Darüber hinaus ist der dblp-Datensatz selbst Untersuchungsgegenstand mehrerer tausend Fachartikel.[10] Insgesamt ist dblp daher für die Informatik sowohl als Recherche-Tool, aber auch als Forschungsdatensatz unverzichtbar geworden.

3.1 *Open Bibliographic Data in Computer Science*

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

For over 20 years now, the “dblp computer science bibliography” has substantially contributed to solving this dilemma in the field of computer science by providing open, quality-checked, and curated bibliographic metadata. The dblp web service supports the computer science research community on several levels, for example by:

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

In addition, the dblp data set itself is object of study of several thousand research articles.[11] Hence, dblp has become indispensable to the computer science community as both a research tool and a research data set.

[10] Google Scholar liefert zum Suchbegriff „dblp“ über 21 700 Treffer; im Einzelnen weisen SpringerLink ca. 2 200 Artikel, Elsevier ScienceDirect über 470 Artikel, die ACM Digital Library ca. 425 Artikel und IEEE Xplore über 190 Artikel nach.

[11] The search term “dblp” results in 21,700 hits at Google Scholar; in particular, SpringerLink lists about 2,200 articles, Elsevier ScienceDirect lists more than 470 articles, the ACM Digital Library lists 425 articles, and IEEE Xplore lists more than 190 articles.

Schloss Dagstuhl und dblp 3.2

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

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

Schloss Dagstuhl and dblp

The cooperation between Schloss Dagstuhl and the dblp computer science bibliography – originally developed at the University of Trier – has existed since late 2010. The commitment of Schloss Dagstuhl to dblp, initially funded by a project of the Leibniz Competition, has been funded directly by Schloss Dagstuhl since June 2013. Since November 2010, Schloss Dagstuhl's dblp team has also been supported by a generous donation from the Klaus Tschira Foundation. Schloss Dagstuhl's own dblp web service at dblp.dagstuhl.de was established in 2012 and complements the dblp service available at the University of Trier at dblp.uni-trier.de. In late 2013, the cooperation agreement between Schloss Dagstuhl and the University of Trier was renewed for another three years.

As part of the consolidation of this cooperation, two and a half Schloss Dagstuhl scientific staff positions – assigned full-time to the support and development of dblp – were created. The dblp advisory board (c.f. Figure 3.1), established in November 2011 at Schloss Dagstuhl, provides scientific supervision and supports dblp with its expertise.

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

Fig. 3.1
dblp Advisory Board.

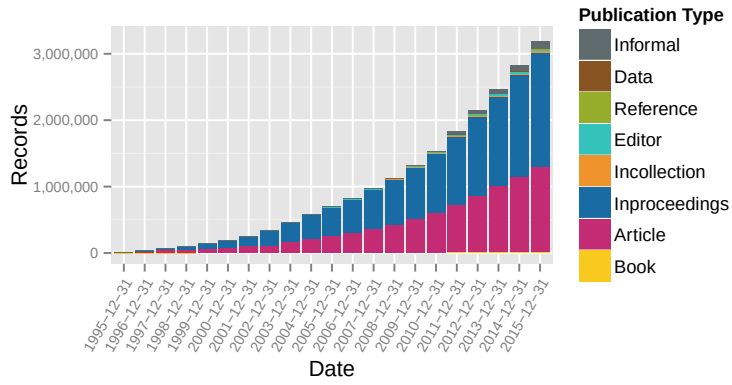

(a) Total number of records by year and type

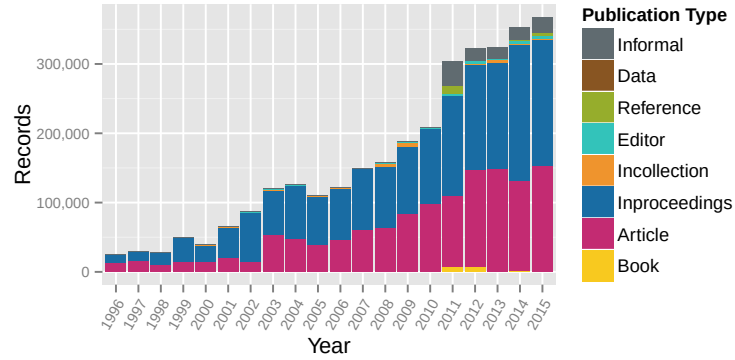

(b) New records by year and type

Fig. 3.2
Development of the dblp data stock.

Statistiken der Datenakquise 3.3 *Data Acquisition Statistics*

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

Zum 31. Dezember 2015 indexierte dblp insgesamt 3 194 657 Publikationen. In dem Zeitraum von Anfang Januar 2015 bis Ende Dezember 2015 wurden dabei 367 864 neue Publikationseinträge in dblp aufgenommen. Diese Aufnahmequote stellt eine Steigerung um weitere 4,1 Prozentpunkte gegenüber der erst im Vorjahr erzielten Rekordaufnahmequote dar. Die neu aufgenommenen Einträge verteilen sich zu 49,5% auf Konferenzbeiträge, zu 41,3% auf Journalartikel, sowie zu 9,2% auf andere Publikationstypen.

Die Anzahl der in dblp erfassten Publikationen überstieg am 18. Juni 2015 die 3-Millionen-Marke. Ein Überblick über die Entwicklung der Datenakquise kann Fig. 3.2a und Fig. 3.2b entnommen werden.

The dblp computer science bibliography indexes conferences and journals on a per-volume basis. Using dblp's own web harvesting software, bibliographic metadata of journal or proceedings volumes are extracted from the publisher's website. This metadata is diligently checked and corrected by the dblp team. The data-cleaning process is assisted by algorithms, but is executed almost exclusively by hand.

Between January 1, 2015, and December 31, 2015, the dblp database grew by 367,864 publication records to reach a total of 3,194,657 records. This number of new records is an 4.1 percent increase when compared to the already high rate of new inclusions reached in 2014. Of these new records, 49.5% have been conference papers, 41.3% have been journal articles, and 9.2% have been other publications.

The total number of publications indexed in dblp passed the 3,000,000 publications milestone on June 18, 2015. The development of the dblp dataset is summarized in Figure 3.2a and Figure 3.2b.

	Trier 1	Trier 2	Dagstuhl
user sessions (visits) per day	28 327	662	510
page views per day	452 089	8 839	14 868
page views per user session	15,9	13,3	29,1
distinct users (IPs) per month	416 413	11 474	7 241
data served per month	861,8 GB	22,0 GB	75,8 GB
as above, including bots	2 206,6 GB	69,5 GB	172,3 GB

Fig. 3.3
Average usage of the three dblp servers in 2015. Trier 1 = http://dblp.uni-trier.de, Trier 2 = http://dblp2.uni-trier.de, Dagstuhl = http://dblp.dagstuhl.de

Nutzungsstatistiken

3.4

Usage Statistics

2015 wurden vom dblp-Team drei offizielle dblp-Server geführt. Die Daten dieser Server werden täglich aktualisiert und miteinander synchronisiert:

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

Seit Mitte 2014 stehen nun auch vergleichbare Nutzerstatistiken von allen drei dblp-Servern zur Verfügung. Dabei ist zu beachten, dass Server Trier 1 aufgrund seiner prominenten Sichtbarkeit in den Google-Suchergebnissen die mit Abstand bekannteste Adresse besitzt.

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

Die angegebenen Daten beinhalten ebenfalls nicht den Datenverkehr, welcher auf die unter der Domain dblp.org betriebenen CompleteSearch-Suchmaschine entfielen. Diese Domain wurde 2015 noch alleine von der Arbeitsgruppe von Prof. Hannah Bast (Universität Freiburg) betrieben und in der ersten Jahreshälfte 2016 vollständig in die dblp-Mirror-Struktur integriert.

In 2015, three official dblp web servers were updated and synchronized on a daily basis:

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

Starting in mid-2014, usage data have been collected on all three mirror sites. The three servers do show a very different rate of usage, with Trier 1 being the by far most widely known server. This is of course due to the fact that server Trier 1 is ranked so highly by the Google search engine.

Figure 3.3 shows the average usage of all three servers in 2015. These figures ignore the traffic caused by known bots and crawlers.

Please note that these figures do not include the traffic of the up to 2015 independently operated CompleteSearch front-end at dblp.org, which was maintained by the research group of Prof. Hannah Bast at the University of Freiburg. Since early 2016, the dblp.org domain is integrated into the dblp mirror structure.

Gemeinsames Projekt von dblp, Zentralblatt MATH und HITS

3.5

Joint Project of dblp, Zentralblatt MATH, and HITS

Die Urheberschaft wissenschaftlicher Publikationen eindeutig zu erkennen und zuzuordnen ist eine der großen Herausforderungen bibliographischer Datendienste. Die Forschung kennt dieses Problem in seiner allgemeinen Form als das Problem der „Entity-Resolution“ oder der „Autorennamen-Disambiguierung“, welches ein wichtiges Forschungsthema im Bereich der linguistischen Datenverarbeitung darstellt. In einem gemeinsamen Projekt wollen sich die Bibliographiedatenbank dblp, das Zentralblatt MATH des FIZ Karlsruhe und das Heidelberger Institut für Theoretische Studien (HITS) diesem Problem annehmen und mit Hilfe des aktuellen Forschungsstandes gemeinsame Lösungsstrategien entwickeln. Die Datensätze von Zentralblatt MATH und dblp teilen dabei die Probleme bei der Identifikation von Autorennamen. Die Kombination beider Datensätze, bestehend aus teils überlappenden und

The correct attribution of scholarly material to their unambiguous authors ranks among the most critical challenges for digital libraries. More generally, the problem of determining which records in a database refer to the same entities is known as “entity resolution” or “author name disambiguation” and constitutes an important field of research within the discipline of natural language processing. In a joint project, the dblp computer science bibliography and the Zentralblatt MATH (located at FIZ Karlsruhe) aim to begin partnering with the Heidelberg Institute for Theoretical Studies (HITS) to find and implement new and state-of-the-art strategies to overcome the challenges of author identification and disambiguation. Zentralblatt MATH and dblp share the challenges associated with author name disambiguation. Due to their partially overlapping, but also partially disjointed data, a joint effort

teils disjunkten Einträgen, stellt dabei eine interessante Möglichkeit dar, Fehler in den Datensätzen aufzudecken und voneinander zu lernen. Die Natural-Language-Processing (NLP) Forschungsgruppe des HITS um Prof. Michael Strube bringt dabei ihre Erfahrung mit graph- und netzwerkbasierten NLP-Methoden bei der Co-Referenz-Resolution und der Konzept- bzw. Entitäts-Disambiguierung ein.

Im Frühjahr 2014 wurde ein Projektantrag für den Leibniz Wettbewerb in der Förderline „Nationale und internationale Vernetzung“ eingereicht. Die beantragte dreijährige Förderung wurde Ende 2014 vom Senat der Leibniz-Gemeinschaft in vollem Umfang bewilligt. Seit dem formalen Projektstart (Anfang Juni 2015) konnte ein weiterer wissenschaftlicher Mitarbeiter zur Verstärkung des dblp-Teams gewonnen werden. In der aktuellen Projektphase wurden Datenaustauschstandards definiert und ein steter Datenaustausch zwischen den Projektpartnern initiiert. Zahlreiche in dblp und zbMATH gemeinsam vertretene Autoren konnten identifiziert und deren Autorenprofile in beiden Datenbanken verlinkt werden. Dies wird eine ganzheitliche Analyse auf beiden Datensätzen ermöglichen und somit auch über die Grenzen der jeweiligen Datensätze vorhandene kontextuelle Zusammenhänge sichtbar machen. Zudem ist derzeit die Erstellung eines hochqualitativen Test- und Trainingsdatensatzes (sogenannte „Goldstandard-Daten“) für die Autorennamen-Disambiguierung in Arbeit. Die nächsten Schritte werden die Etablierung von Evaluations- und Testumgebungen sowie die Adaption von NLP-Methoden für die Entitäts-Disambiguierung beinhalten. Das Projekt läuft bis Juni 2018.

to identify authors based on the combination of the two data sets appears to be very promising. The Natural Language Processing (NLP) Group at the HITS, lead by Prof. Michael Strube, joins the project by providing its extensive experience with graph-based and network methods for NLP tasks such as co-reference resolution, cross-document co-reference resolution, concept and entity disambiguation.

In early 2014, a project proposal was submitted to the “National and International Networking” funding line of the Leibniz Competition. The project has been picked up for funding and has been formally started in July 2015. Since then, one further scientific project staff member reinforced the dblp team. In its current phase, the project partners have exchanged first test data-sets and developed a common data-exchange standard. Numerous common author identities have already been identified in and linked between the dblp and zbMATH data sets. This will allow the joint data of both databases to be analyzed as a whole and make new contextual information visible. Currently, the construction of a high quality training and test data set (“gold standard data”) for the author disambiguation task is in progress. Next steps will include the specification of evaluation metrics and the adaptation of state-of-the-art NLP learning techniques to the problem-specific domain using the gold standard data. The project will conclude in June 2018.

4 Dagstuhl Publishing
Dagstuhl Publishing

Dagstuhl Publishing

Portfolio

4.1

Portfolio

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

The scientific community appreciates the Open Access publishing services offered by Schloss Dagstuhl. The portfolio covers series related to events at Schloss Dagstuhl (*Dagstuhl Reports*, *Dagstuhl Manifestos*, *Dagstuhl Follow-Ups*) and series for conferences and workshops held outside of Schloss Dagstuhl (*OASIcs* and *LIPIcs*). The scholarly journal *LITES* has been running since 2013, when it was launched. In 2015, the DARTS series was established which aims at publishing research artefacts.

Dagstuhl Reports

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

Die Zeitschrift wurde 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Seminaren und Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 12.4) 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 2015 wurde für 71 Seminare und Perspektiven-Workshops ein Bericht in der Reihe *Dagstuhl Reports* veröffentlicht. An dieser Stelle bedanken wir uns ganz herzlich bei den Organisatoren und Kollektoren für die erfolgreiche Zusammenarbeit.

Dagstuhl Reports

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

The periodical started with the first seminars of January 2011 and publishes in monthly issues reports on seminars and workshops that took place on a given month. The content is not peer-reviewed. The Scientific Directorate (see Fig. 12.4) 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).

71 Dagstuhl Seminars and Dagstuhl Perspectives Workshops that took place in 2015 have published a report. We would like to take this opportunity to cordially thank all organizers and collectors for their successful collaboration.

Dagstuhl Manifestos

Seit 2011 werden in der Zeitschrift *Dagstuhl Manifestos*[13] die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 12.4) fungiert hier ebenfalls als Herausgebergremium. Die Ausgabe für 2015 enthält zwei Manifestos, siehe Fig. 4.1.

Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*[13] in Open Access manner. The Scientific Directorate (see Fig. 12.4) acts as the editorial board of the journal. The 2015 volume includes two Dagstuhl Manifestos; see Fig. 4.1.

Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*[14] ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop basiert. Für jedes Buch ist ein Antrag notwendig, der vom wissenschaftlichen Direktorium (welches als Herausgebergremium verantwortlich ist) begutachtet

Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups*[14] book series is devoted to peer-reviewed collections of original research works that are rooted in a dedicated Dagstuhl Seminar or Dagstuhl Perspectives Workshop. Each book needs a proposal, which is reviewed and finally approved by the Scientific Directorate (which is in charge as editorial board). In 2015, no volume

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

Dagstuhl Publishing

und freigegeben werden muss. In 2015 wurde kein Buch in der Reihe veröffentlicht, jedoch gab es einen Antrag ausgehend vom Dagstuhl Seminar 15301 „The Constraint Satisfaction Problem: Complexity and Approximability“, welcher angenommen wurde. Dieser Band soll im Spätsommer 2016 erscheinen.

OASIcs: OpenAccess Series in Informatics

Die *OASIcs*-Reihe[15] 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 2015 wurden 6 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 4.3.

LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe[16] veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausgebergremium besteht aus einschlägig bekannten Wissenschaftlern und wurde bis Mai 2015 von Pascal Weil als Hauptherausgeber geleitet. Nach Auslauf seiner Amtszeit wurde mit Wolfgang Thomas ein sehr erfahrener und umsichtiger Hauptherausgeber gewählt, der nun bis Mai 2017 im Amt sein wird.

was published in the series. However, a proposal based on Dagstuhl Seminar 15301 “The Constraint Satisfaction Problem: Complexity and Approximability” was submitted and accepted. This volume is scheduled for late summer 2016 for publishing.

OASIcs: OpenAccess Series in Informatics

The *OASIcs* series[15] 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 2015, Dagstuhl published 6 *OASIcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 4.3.

LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series[16] publishes proceedings of leading conferences in the area of informatics. An international editorial board of renowned researchers supervises the conferences that are accepted for LIPIcs and was headed until May 2015 by Pascal Weil. After his term ended, Wolfgang Thomas – who is an experienced and far-seeing researcher – was elected as the new chair until May 2017.

In 2015, the terms of Susanne Albers, Michael Mitzenmacher and Madhavan Mukund were extended until May

[14] http://drops.dagstuhl.de/dfu
[15] http://drops.dagstuhl.de/oasics

Connecting Performance Analysis and Visualization http://dx.doi.org/10.4230/DagMan.5.1.1 based on Dagstuhl Perspectives Workshop 14022 14022
Privacy and Security in an Age of Surveillance http://dx.doi.org/10.4230/DagMan.5.1.25 based on Dagstuhl Perspectives Workshop 14401

Fig. 4.1
Manifestos published in the 2015 volume of the journal *Dagstuhl Manifestos*.

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

Fig. 4.2
OASIcs Editorial Board.

Die Amtszeiten von Susanne Albers, Michael Mitzenmacher und Madhavan Mukund wurden 2015 nach einem anonymen Wahlverfahren innerhalb des Herausgebergremiums bis Mai 2019 verlängert. Siehe auch Fig. 4.4.

In 2015 wurden Tagungsbände von 16 Konferenzen veröffentlicht, so viel wie noch nie zuvor; siehe Fig. 4.5.

Der TYPES-Workshop wurde vom Herausgebergremium im Rahmen der bestehenden Kooperation erneut evaluiert und für weitere fünf Jahre (2015–2019) aufgenommen.

In 2015 gab so erneut viele Bewerbungen bei LIPIcs und setzte so die große Nachfrage aus dem Vorjahr fort. Die große Anzahl an Bewerbungen sind die erfreulichen Ergebnisse unserer langjährigen Bemühungen, einige der wichtigsten Konferenzen an LIPIcs zu binden. In Fig. 4.6 sind alle Konferenzen aufgelistet, deren Anträge bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischweise 5 Jahre) eingegangen wurde.

2019, based on an anonymous voting within the editorial board. See also Fig. 4.4.

The series published the proceedings of 16 major conferences in 2015, marking a record high since the series was started; see Fig. 4.5.

The TYPES workshop was re-evaluated by the LIPIcs editorial board and accepted for another five-year period (2016–2019).

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2015 has seen several applications for LIPIcs, continuing the high interest from the previous year. Fig. 4.6 lists all conferences that have been accepted for a cooperation covering several years (typically 5 years).

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

Vol. 44 | SynCoP'15 | 2nd International Workshop on Synthesis of Complex Parameters http://www.dagstuhl.de/dagpub/978-3-939897-82-8
Vol. 45 | CMN'15 | 6th Workshop on Computational Models of Narrative http://www.dagstuhl.de/dagpub/978-3-939897-93-4
Vol. 46 | WPTE'15 | 2nd International Workshop on Rewriting Techniques for Program Transformations and Evaluation http://www.dagstuhl.de/dagpub/978-3-939897-94-1
Vol. 47 | WCET'15 | 15th International Workshop on Worst-Case Execution Time Analysis http://www.dagstuhl.de/dagpub/978-3-939897-95-8
Vol. 48 | AMTOS'15 | 15th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems http://www.dagstuhl.de/dagpub/978-3-939897-99-6
Vol. 49 | ICCSW'15 | 2015 Imperial College Computing Student Workshop http://www.dagstuhl.de/dagpub/978-3-95977-000-2

Fig. 4.3
OASIcs volumes published in 2015.

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

Fig. 4.4
LIPIcs Editorial Board.

Vol. 30 | STACS'15 | 32nd International Symposium on Theoretical Aspects of Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-78-1
Vol. 31 | ICDT'15 | 18th International Conference on Database Theory http://www.dagstuhl.de/dagpub/978-3-939897-79-8
Vol. 32 | SNAPL'15 | 1st Summit on Advances in Programming Languages http://www.dagstuhl.de/dagpub/978-3-939897-80-4
Vol. 33 | CCC'15 | 30th Conference on Computational Complexity http://www.dagstuhl.de/dagpub/978-3-939897-81-1
Vol. 34 | SoCG'15 | 31st International Symposium on Computational Geometry http://www.dagstuhl.de/dagpub/978-3-939897-83-5
Vol. 35 | CALCO'15 | 6th Conference on Algebra and Coalgebra in Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-84-2
Vol. 36 | RTA'15 | 26th International Conference on Rewriting Techniques and Applications http://www.dagstuhl.de/dagpub/978-3-939897-85-9
Vol. 37 | ECOOP'15 | 29th European Conference on Object-Oriented Programming http://www.dagstuhl.de/dagpub/978-3-939897-86-6
Vol. 38 | TLCA'15 | 13th International Conference on Typed Lambda Calculi and Applications http://www.dagstuhl.de/dagpub/978-3-939897-87-3
Vol. 39 | TYPES'14 | 20th International Conference on Types for Proofs and Programs http://www.dagstuhl.de/dagpub/978-3-939897-88-0
Vol. 40 | APPROX/RANDOM'15 | Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques http://www.dagstuhl.de/dagpub/978-3-939897-89-7
Vol. 41 | CSL'15 | 24th EACSL Annual Conference on Computer Science Logic http://www.dagstuhl.de/dagpub/978-3-939897-90-3
Vol. 42 | CONCUR'15 | 26th International Conference on Concurrency Theory http://www.dagstuhl.de/dagpub/978-3-939897-91-0
Vol. 43 | IPEC'15 | 10th International Symposium on Parameterized and Exact Computation http://www.dagstuhl.de/dagpub/978-3-939897-92-7
Vol. 44 | TQC'15 | 10th Conference on the Theory of Quantum Computation, Communication and Cryptography http://www.dagstuhl.de/dagpub/978-3-939897-96-5
Vol. 45 | FSTTCS'15 | 35th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-97-2

Fig. 4.5
LIPIcs volumes published in 2015.

CPM	Annual Symposium on Combinatorial Pattern Matching accepted for 2016–2020
ECOOP	European Conference on Object-Oriented Programming accepted for 2015–2019
FSCD	Formal Structures for Computation and Deduction accepted for 2016–20120
FUN	International Conference on Fun with Algorithms accepted for 2016–2020
ISAAC	International Symposium on on Algorithms and Computation accepted for 2016–2020
MFCS	Mathematical Foundations of Computer Science accepted for 2016–2020
OPODIS	International Conference on Principles of Distributed Systems accepted for 2015–2019
SWAT	Scandinavian Symposium and Workshops on Algorithm Theory accepted for 2016–2020

Fig. 4.6
Conferences accepted in 2015 for publication in LIPIcs.

Prof. Alan Burns, DPhil University of York, UK | Editor-in-Chief
Prof. Sang Lyul Min, Ph. D. Seoul National University, South Korea | Subject area: Architecture, platforms
Prof. Dr. Marco di Natale Scuola Superiore Santa Anna, Italy | Subject area: Automotive applications
Dr. Virginie Wiels ONERA, France | Subject area: Avionics applications
Prof. Karl-Erik Arzen, Ph. D. Lund University, Sweden | Subject area: Control
Prof. Steve Goddard, Ph. D. University of Nebraska-Lincoln, US | Subject area: Cyber-physical systems
Prof. Dr. Axel Jantsch Royal Institute of Technology Stockholm, Sweden | Subject area: Distributed embedded systems and networks
Prof. Bashir Al Hashimi University of Southampton, UK | Subject area: Energy-efficiency
Prof. Mateo Valero, Ph. D. Technical University of Catalonia | Subject area: High-performance embedded systems
Prof. Dr. Martin Fränzle Carl von Ossietzky University Oldenburg, Germany | Subject area: Hybrid systems
Prof. Dr. Samarjit Chakraborty Technical University Munich, Germany | Subject area: Multimedia applications
Prof. Dr. Gernot Heiser University of New South Wales, Australia | Subject area: Operating systems
Prof. Dr. Lothar Thiele ETH Zürich, Switzerland | Subject area: Performance and wireless sensor networks
Dr. Neil Audsley University of York, UK | Subject area: Real time
Prof. Sanjoy Baruah, Ph. D. University of North Carolina at Chapel Hill, US | Subject area: Scheduling
Prof. Dr. Florence Maraninchi University of Grenoble, France | Verimag Lab, France | Subject area: Verification, formal methods, model-based design

Fig. 4.7
LITES Editorial Board.

Dagstuhl Publishing

LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*[17] veröffentlicht begutachtete Beiträge zu allen Aspekten eingebetteter Systeme. In 2012 wurde die Zeitschrift gegründet und in 2013 wurde der Betrieb aufgenommen. Ein breit aufgestelltes Team an erfahrenen Wissenschaftlern, die sich für ihr jeweiliges Fachgebiet verantwortlich zeichnen (siehe Fig. 4.7), begutachtet alle eingereichten Arbeiten.

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. Die APC von 100 € ist momentan für den Zeitraum 2013–2015 sichergestellt Dank finanzieller Unterstützung von Google und der Klaus Tschira Stiftung.

Nachdem in 2014 eine Absichtserklärung über die gemeinsame Herausgeberschaft mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*[18] der Fachgesellschaft *European Design and Automation Association (EDAA)*[19] unterschrieben wurde, konnte nun in 2015 die Zusammenarbeit durch einen formalen Kooperationsvertrag fixiert werden. Die Fachgruppe für die Besetzung des Herausgebergremiums verantwortlich, während Schloss Dagstuhl die administrativen Aufgaben der Herausgeberschaft übernimmt. Unabhängig von der Kooperation verbleibt die Zeitschrift in Besitz von Schloss Dagstuhl.

In 2015 wurden zwei Ausgaben von *LITES* mit insgesamt 3 Artikeln veröffentlicht.

DARTS: Dagstuhl Artifacts Series

In der Reihe *DARTS*[20] werden qualitätsgesicherte Forschungsdaten und -artefakte veröffentlicht. Die Reihe hat dabei die Struktur einer Zeitschrift. In 2015 wurde die erste Ausgabe mit 12 Artefakten veröffentlicht. Diese Artefakte wurden im Rahmen der 29. European Conference on Object-Oriented Programming (ECOOP'15) – deren Konferenzband als Volume 37 in LIPIcs veröffentlicht wurde – evaluiert, wobei das Regelwerk „Artifact Evaluation for Software Conferences"[21] angewendet wurde. Jedes Artefakt wird mit einer separaten Beschreibung veröffentlicht.

Die Veröffentlichung und Bereitstellung von Forschungdaten und -artefakten ist aktuell ein wichtiges Thema in den wissenschaftlichen Disziplinen und bei den Forschungsfördereinrichtungen. Im Bereich der Informatik wird dieses Theme ebenfalls diskutiert. In 2015 gab es zum Beispiel einen Perspektiven-Workshop mit dem Titel „Artifact Evaluation for Publications"[22], der in 2016 durch zwei weitere Seminare ergänzt wird: „Reproducibility of Data-Oriented Experiments in e-Science"[23] and „Rethinking Experimental Methods in Computing"[24].

Schloss Dagstuhl unterstützt mit DARTS die Wissenschaftsgemeinde in der Informatik bei dem Wunsch,

LITES: Leibniz Transactions on Embedded Systems

The *LITES*[17] journal publishes original peer-reviewed articles on all aspects of embedded computer systems via Open Access. The journal was established in 2012 and started operating in early 2013. A broad team of experienced researchers, acting as editorial board (see Fig. 4.7), reviews all submitted contributions.

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. The APC of 100 € is guaranteed for the 2013–2015 period thanks to support from sponsors like Google and the Klaus Tschira Stiftung.

After a Memorandum of Understanding between Schloss Dagstuhl and the *EMbedded Systems Special Interest Group (EMSIG)*[18] of the *European Design and Automation Association (EDAA)*[19] regarding the joint publication of LITES was signed in 2014, the formal cooperation contract was now signed in 2015. The special interest group is responsible for appointing the editorial board, while Schloss Dagstuhl takes over the administrative tasks of the publication. Independently of this cooperation, Schloss Dagstuhl retains ownership of the journal.

In 2015, two issues of *LITES* containing 3 articles in total were published.

DARTS: Dagstuhl Artifacts Series

The *DARTS* series[20] publishes evaluated research data and artifacts. It is organized as a periodical. In 2015, one volume containing one issue with 12 artifacts was published. This first issue contains the research artifacts of the 29th European Conference on Object-Oriented Programming (ECOOP'15), whose conference proceedings were published as volume 37 in the LIPIcs series. Each artifact is published with a separate description and was evaluated according to the guidelines for "Artifact Evaluation for Software Conferences"[21].

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

With DARTS, Schloss Dagstuhl is aiming to support the computing research community with a publishing venue dedicated to research data and artifacts. Especially,

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

Dagstuhl Publishing

Forschungsdaten und -artefakte in einer geeigneten Reihe zu veröffentlichen. Hierbei berücksichtigt DARTS insbesondere auch die Publikationskultur in der Informatik mit ihrem Schwerpunkt auf Konferenzbandveröffentlichungen.

DARTS takes into account the publication culture in computer science which focusses on conference proceedings publications.

Infrastruktur 4.2 *Infrastructure*

Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 4.8. Die Bände aus der Reihe *LIPIcs* werden beim Conference Proceedings Citation Index (CPCI), welcher vom Medienkonzern Thomson Reuters unterhalten wird, eingereicht; zudem werden diese seitens SCOPUS in deren Katalog aufgenommen. Die Reihen *LIPIcs* und *OASIcs* sowie die Zeitschrift *LITES* sind zudem im Directory of Open Access Journals (DOAJ) gelistet, siehe Fig. 4.8.

Die Bände der LIPIcs-Reihe werden bei Scopus [25] regelmäßig indiziert. Zudem unterstützen die technischen Schnittstellen die Datenakquisition durch GoogleScholar, so das die Publikationen sichtbarer und besser recherchierbar sind.

Indexing

All series of the publication portfolio are listed in *dblp*; see Fig. 4.8. The LIPIcs volumes are submitted to the Conference Proceedings Citation Index (CPCI), maintained by the Thomson Reuters media group; additionally, SCOPUS is integrating them into their catalog. The *LIPIcs* and *OASIcs* series as well as the journal *LITES* are also listed in the Directory of Open Access Journals (DOAJ), see Fig. 4.8.

LIPIcs volumes are regularly indexed in Scopus[25] and the technical interface of our publication server enables harvesting according to the guidelines of GoogleScholar. GoogleScholar regularly retrieves metadata and full-texts from our server.

LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*[26] 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.

LeibnizOpen

The Leibniz Association has established the *LeibnizOpen*[26] 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.

AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Sichtbarkeit und Strategie“[27] mit organisiert, welcher bereits der dritte Workshop in Folge war seit 2013. Der Workshop fand am 22. und 23. Januar 2015 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt. Beim Workshop haben 36 Teilnehmern aus den Verlagsabteilungen von ungefähr 20 Leibniz-Instituten teilgenommen. Für 2017 ist eine Nachfolgeveranstaltung in Planung.

Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Visibility and Strategy”[27] was initiated and coordinated as part of our membership in the Open Access working group of the Leibniz Association. The workshop took place at the Leibniz Association headquarters in Berlin on January 22 and 23, 2015 and brought together 36 professionals in charge of publishing activities at about 20 Leibniz institutes. For 2017 a follow-up workshop is planned.

[20] http://www.dagstuhl.de/darts
[21] http://www.artifact-eval.org/
[22] ttp://www.dagstuhl.de/15452
[23] http://www.dagstuhl.de/16041
[24] http://www.dagstuhl.de/16111
[25] http://www.scopus.com
[26] http://www.leibnizopen.de/
[27] https://www.dagstuhl.de/dagpub/journalmanagement-leibniz/2015-01-22-workshop/

Dagstuhl Publishing

AG Open Access der Schwerpunktinitiative „Digitale Information“

Die Allianz der deutschen Wissenschaftsorganisationen, zu der neben der Max-Planck-Gesellschaft, der Helmholtz-Gemeinschaft, sowie weiteren Organisationen auch die Leibniz-Gemeinschaft gehört, hat eine Schwerpunktinitiative „Digitale Information“ ins Leben gerufen, bei der auch das Thema *Open Access* als Handlungsfeld vertreten ist. Mit Dr. Marc Herbstritt wurde seitens der Leibniz-Gemeinschaft ab Juli 2013 ein Mitglied des wissenschaftlichen Stabs von Schloss Dagstuhl in die Arbeitsgruppe „Open Access“[28] berufen.

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

Technisches Back-end: DROPS

Über den Dagstuhl Research Online Publication Server (DROPS)[29] werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative[30] berücksichtigt, wodurch alle nötigen Metadaten zu jeder Publikation gespeichert werden

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

The Alliance of German Science Organizations, to which – among others – the Max Planck Society, the Helmholtz Association and also the Leibniz Association belong, has established a priority initiative “Digital Information” where *Open Access* is handled as a core activity. Since July 2013, Dagstuhl scientific staff member Dr. Marc Herbstritt has collaborated with this working group as the delegated representative of the Leibniz Association.[28]

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

Back-end: DROPS

All items published by the center are administered via the Dagstuhl Research Online Publication Server (DROPS)[29]. The general guidelines of the Dublin Core initiative[30] applicable to online publications are adhered to, meaning that all the requisite metadata of each publication is stored, thus ensuring availability in the long

[28] http://www.allianzinitiative.de/de/handlungsfelder/open_access/

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

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

Fig. 4.8
Indexing of Dagstuhl Publishing series in dblp and DOAJ.

Dagstuhl Publishing

und die Langzeitverfügbarkeit sichergestellt wird. Die Online-Publikationen sind zitierfähig und stehen einer grossen Leserschaft zur Verfügung. Als technische Grundlage dient eine adaptierte Version des OPUS-Systems.[31]

Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)[32] zur (digitalen) Langzeitarchivierung eingereicht.

Mirroring

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

- io-port.net: Das unter Leitung des FIZ Karlsruhe, Leibniz-Institut für Informationsinfrastruktur, organisierte Informatik-Publikations-Portal io-port.net spiegelt alle Bände der LIPIcs-Reihe.[33] In 2011 wurde die bestehende Verbindung durch eine gemeinsame Kooperationserklärung gefestigt.
- SunSite Central Europe: Der Sun-Server-Park, der an der RWTH Aachen unter Leitung von Prof. Matthias Jarke betrieben wird, bietet eine Heimat für zahlreiche Software-Archive als auch Publikationen. Der gesamte DROPS-Bestand wird nun in regelmäßigen Abständen auf der SunSite Aachen gespiegelt.[34]

term. This enables the online publications to be cited by and accessible to a wide readership. The technical basis for this is an adapted version of the OPUS system.[31]

Long-term Archiving

All publications are submitted to the German National Library (D-NB)[32] for (digital) long-term archiving.

Mirroring

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

- io-port.net: The informatics publication portal organized under the auspices of io-port.net, FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, mirrors all volumes of the LIPIcs series[33]. In 2011, the existing affiliation was consolidated by a memorandum of understanding.
- SunSite Central Europe: The Sun server park, located at the Aachen University of Technology and operated under the guidance of Prof. Matthias Jarke, is home to numerous software archives and publications. All the DROPS assets are now mirrored at regular intervals on the Aachen SunSite.[34]

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

5 Resonanz
Feedback

Resonanz von Seminarteilnehmern

5.1

Feedback from Seminar Participants

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

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

Martin Erwig

13091 – Analysis, Test and Verification in the Presence of Variability | Dagstuhl Seminar | http://www.dagstuhl.de/13091

Here is a short success story that evolved over two Dagstuhl seminars. Having been invited to the 2011 Seminar 11021 (Feature-Oriented Software Development (FOSD)), I met Sven Apel (one of the organizers) and Christian Kästner. This allowed my then-Ph.D. student Eric Walkingshaw and myself to learn about our different approaches to variation representation and led to an exchange of ideas in the time following the seminar. In a later seminar in 2013 that addressed a similar topic (13091 - Analysis, Test and Verification in the Presence of Variability) we also connected with Eric Bodden and forged a collaboration to explore a new research area that is centered around the idea of variational data structures. A first outcome of this collaboration is the following paper that appeared last year. Variational Data Structures: Exploring Tradeoffs in Computing with Variability Eric Walkingshaw, Christian Kästner, Martin Erwig, Sven Apel, and Eric Bodden Onward! 2014 Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software, Pages 213-226 ACM New York, NY, USA ©2014 DOI: 10.1145/2661136.2661143 Following this publication, Eric Walkingshaw, Christian Kästner, and I submitted a joint NSF grant proposal. This proposal is still under review. In case it gets funded I will let you know. This research would not have happened without the Dagstuhl seminars, which have provided a unique setting to develop and explore new research agendas. A big "Thank You"to all the sponsors and supporters of the Dagstuhl seminar series. Their support makes a huge difference and has an important impact on the future of research and society. Sincerely, Martin Erwig

Yann Ollivier

15211 – Computational Geometry | Dagstuhl Seminar | http://www.dagstuhl.de/15211

This was my first experience at Dagstuhl and I must say I am thoroughly impressed and I loved it. I deeply appreciated the efforts to make our stay comfortable and productive, with comfortable and sound-proof rooms, devoted staff, nice catering, tea and coffee available at every time, beautiful surroundings, etc., and especially the general atmosphere of trust.

Rolf Klein

15111 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | http://www.dagstuhl.de/15111

I'm very grateful for having been invited to Dagstuhl several times, as a participant or a co-organizer. In my opinion Dagstuhl is not only the most prominent computer science institution of Germany; it is a unique place world-wide. Counting the number of publications, cooperations, and projects that result from a given seminar or a seminar series may lead to impressive numbers, but I think there is more. That the best and the most promising researchers can regularly spend time together at a place of this special character promotes the development of whole scientific communities. Important ingredients seem to be: the absence of time pressure and bureaucratic hassle, the beautiful environment, and the feeling that all problems of daily life are being taken care of by an extremely competent team. The excellent library, too. Please do not change this great concept!

Resonanz unserer Seminarorganisatoren

5.2 *Feedback from Seminar Organizers*

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

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

Simon M. Lucas

15051 – Artificial and Computational Intelligence in Games: Integration | Dagstuhl Seminar | http://www.dagstuhl.de/15051

Here are some success stories where we believe the previous seminar on Artificial and Computational Intelligence in Games played an important role in bringing people together, brainstorming and developing ideas that subsequently became major activities:

- This EPSRC Funded Doctoral Training Programme in Intelligent Games and Game Intelligence: IGGI-Webpage[35]
- The General Video Game AI Competition: GVGAI-Webpage[36] (an on-going series of AI competitions tied in to various international conferences)
- Journal paper on General Video Game AI[37] (there were also many conference papers associated with this)

[35] http://iggi.org.uk
[36] http://gvgai.net/
[37] http://dx.doi.org/10.1109/TCIAIG.2015.2402393

Organizers of Dagstuhl Seminar 15111

15111 – Computational Geometry | Dagstuhl Seminar | http://www.dagstuhl.de/15111

This year, 41 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive. No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at the Dagstuhl Seminar. We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, Schloss Dagstuhl itself is a great strength of the seminar.

Organizers of Dagstuhl Seminar 15121

15121 – Mixed Criticality on Multicore/Manycore Platforms | Dagstuhl Seminar | http://www.dagstuhl.de/15121

As organizers, we would like to thank Prof. Reinhard Wilhelm for encouraging us to submit the seminar proposal, Dagstuhl's Scientific Directorate for allowing us to run a seminar on mixed criticality systems, and to the staff at Schloss Dagsthul for their superb support during the seminar itself. Finally, we would like to thank all of the participants for their strong interaction, presentations, group discussions, and work on open problems, sometimes into the early hours of the morning. We were very pleased to hear about the progress of new found collaborations, and to receive such positive feedback about the seminar itself. Thank you to everyone who participated for a most enjoyable and fruitful seminar.

Organizers of Dagstuhl Seminar 15021

15021 – Concurrent computing in the many-core era | Dagstuhl Seminar | http://www.dagstuhl.de/15021

As detailed in the rest of this report, the seminar has allowed the community to make significant progress on a number of important questions pertaining to concurrent computing, while at the same time defining a research agenda for the next few years. Participants provided very positive feedback following the seminar and expressed strong interest in follow-up events. Organizers strongly support the continuation of this series of seminars on concurrent computing, one of the most important and challenging fields in the era of multiand many-core systems.

Ute Schmid

15442 – Approaches and Applications of Inductive Programming | Dagstuhl Seminar | http://www.dagstuhl.de/15442

Das Dagstuhl-Format ist meiner Meinung nach die aller produktivste Organisationsform für intensiven wissenschaftlichen Austausch. Nirgends sonst erlebe ich eine so intensive Auseinandersetzung mit einem Thema. Nirgends sonst hat man so viel Zeit, eine Sache wirklich durchzudiskutieren. Nirgends sonst erlebe ich eine so offene Atmosphäre, wo es nicht um das möglichst glatte Präsentieren von Erreichtem sondern um das offene Aufzeigen von offenen Fragen in der eigenen Forschung geht. Nirgends sonst habe ich die Möglichkeit, ganz andere Ansätze innerhalb meines eigenen Fachs und ganz andere Perspektiven von anderen Disziplinen auf ein Thema kennenzulernen. Und nirgends sonst kann man sich durch die tolle Rundumversorgung so intensiv auf den wissenschaftlichen Austausch konzentrieren.

Steve Vickers

15441 – Duality in Computer Science | Dagstuhl Seminar | http://www.dagstuhl.de/15441

I didn't get round to filling in the survey, but I fully agree with the positive tone of the responses. It was a very enjoyable meeting, with an excellent mix of participants and good opportunities for new directions and collaborations. All this, of course, was enhanced by the special atmosphere of Dagstuhl.

Resonanz in Sozialen Netzwerken

5.3 ***Feedback in Social Media***

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

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

Annegret Kramp-Karrenbauer

15503 – Lehrerfortbildung in Informatik | Educational Event | https://twitter.com/_A_K_K_/status/674614330739048448

A. Kramp-Karrenbauer @_A_K_K_ (4:39 PM – 9 Dec 2015)
Herzlichen Glückwunsch zu 25 Jahren erfolgreicher Lehrerfortbildung auf Schloss @dagstuhl #Informatik

@hshdinformatik

15315 – Fakultätsklausur der Fakultät für Informatik – SRH Hochschule Heidelberg | Research Group Meeting | https://twitter.com/InfoMediaDesign/status/626087186300071937

Fakultät Informatik @hshdinformatik (7:49 PM – 28 Jul 2015)
Wir sind zurück von Schloss @dagstuhl - Ein toller Ort um die wichtigen Entscheidungen zu treffen, die unsere Zukunft prägen.

Andrea Cerone

15191 – Compositional Verification Methods for Next-Generation Concurrency | Dagstuhl Seminar | https://twitter.com/acerone85/status/593024075204792320

@acerone85 (2:08 PM – 28 Apr 2015)
Very happy to be participating at the forthcoming @dagstuhl seminar „Compositional Verification Methods for Next-Generation Concurrency“

tom vogel

14382 – Control Theory meets Software Engineering | GI-Dagstuhl Seminar | https://twitter.com/tomvog/status/568343157315903488

tom vogel @tomvog (10:35 AM – 19 Feb 2015)
Paper „Software Engineering Meets Control Theory“, joint outcome of the GI @dagstuhl seminar 14382 participants, accepted at #SEAMS15

Alex J. Champandard

15051 – Artificial and Computational Intelligence in Games: Integration | Dagstuhl Seminar | https://twitter.com/alexjc/status/562203715520966656

Alex J. Champandard @alexjc (11:59 AM – 2 Feb 2015)
Recovered from an intense week at @dagstuhl. Thinking about the major trends and changes in the field over the past year...

Michael Cook

15051 – Artificial and Computational Intelligence in Games: Integration | Dagstuhl Seminar | https://twitter.com/mtrc/status/561172189257687040

Michael Cook @mtrc (3:40 PM – 30 Jan 2015) On my way home after a week at @dagstuhl that is honestly hard to describe. I feel like a year of work just got done. Report later!

Resonanz im Fragebogen 5.4 *Seminar Survey Feedback*

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

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

15041 – Model-driven Algorithms and Architectures for Self-Aware Computing Systems | Dagstuhl Seminar | http://www.dagstuhl.de/15041

Enthusiastic participants and organisers with lot of energy and flexibility

15052 – Empirical Evaluation for Graph Drawing | Dagstuhl Seminar | http://www.dagstuhl.de/15052

The new idea with having trainers from other disciplines was very, very useful.

15062 – Domain-Specific Languages | Dagstuhl Seminar | http://www.dagstuhl.de/15062

The seminar was very well organized. The poster session was a very good way of introducing people.

15062 – Domain-Specific Languages | Dagstuhl Seminar | http://www.dagstuhl.de/15062

Plenty of time for discussions. Enough structure to force us to do stuff, but not so much that it felt like a workshop. Plenty of free-form activity time. At first I didn't like the fact that we weren't told the program until the night before: I didn't know when I would be giving a talk, or what would be happening. Then I figured out the organizers were intentionally disrupting our behaviour. At that point, I got into the flow of it. I can't think of anything bad – maybe more opportunities for demos? A bit more structure for the first couple of talks???

15061 – Non-Zero-Sum-Games and Control | Dagstuhl Seminar | http://www.dagstuhl.de/15061

A place dedicated to efficient work and ideas sharing.

15081 – Holistic Scene Understanding | Dagstuhl Seminar | http://www.dagstuhl.de/15081

Best: atmosphere allows for discussions with other researchers and the possibility to approach people any time during the seminar. Worst: none.

15081 – Holistic Scene Understanding | Dagstuhl Seminar | http://www.dagstuhl.de/15081

The relatively small number of participants results in fruitful and lively discussions.

15111 – Computational Geometry | Dagstuhl Seminar | http://www.dagstuhl.de/15111

I really liked the mix of topics (in particular topological data analysis) and the mix of younger participants. I felt like I really learned a lot from the talks that I can immediately apply to my research.

15161 – Advanced Stencil-Code Engineering | Dagstuhl Seminar | http://www.dagstuhl.de/15161

The seminar was all round great – excellent organization and excellent format; ideal breaks between talks. The best aspect was the daily end of day discussions.

15182 – Qualification of Formal Methods Tools | Dagstuhl Seminar | http://www.dagstuhl.de/15182

As an academic, I learned a lot about practical issues in several areas of industry, and also about recent or ongoing work that I was not aware of.

15171 – Theory and Practice of SAT Solving | Dagstuhl Seminar | http://www.dagstuhl.de/15171

The best aspect was the mix of top researchers from very different areas. The talks were generally of very high quality.

15192 – The Message in the Shadow: noise or knowledge? | Dagstuhl Seminar | http://www.dagstuhl.de/15192

A seminar group with experts from a large variety of different areas who communicated extremely well and constructively.

15192 – The Message in the Shadow: noise or knowledge? | Dagstuhl Seminar | http://www.dagstuhl.de/15192

Excellent mix of people, positive attitude of participants; atmosphere created by organizers.

15211 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | http://www.dagstuhl.de/15211

Best: Having all important people of the field in one place for a whole week, in a non-hectic atmosphere that cannot be found at a conference Worst: There was nothing to complain

15211 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | http://www.dagstuhl.de/15211

Best: exchanging ideas emerging from excellent presentations with many different open minds. Discussions with people that I wouldn't otherwise have a chance to talk to. The "Theory of Evolutionary Algorithms" seminar series is unbelievably inspiring – I'm leaving with many new ideas for research, new contacts, and new collaborations. I've been looking forward to this all year and my expectations were exceeded. Worst: the seminar only takes place every other year!

15351 – Computational Mass Spectrometry | Dagstuhl Seminar | http://www.dagstuhl.de/15351

It's a "scientific holiday" in the best sense. Highly stimulating, but also relaxing. Much deeper interaction with other participants than at a typical conference. The environment and Dagstuhl setup are superb.

15351 – Computational Mass Spectrometry | Dagstuhl Seminar | http://www.dagstuhl.de/15351

Though I felt the seminar to be arguably grounded in computational proteomics, the inclusion of experts more directly aligned with metabolomics and statistics was one of the best aspects of this seminar.

15371 – Quantum Cryptanalysis | Dagstuhl Seminar | http://www.dagstuhl.de/15371

It was great that there were people from quantum algorithms and quantum information theory on the one hand and cryptographers on the other hand. These groups are otherwise quite separated – sadly.

15392 – Measuring the Complexity of Computational Content: Weihrauch Reducibility and Reverse Analysis | Dagstuhl Seminar | http://www.dagstuhl.de/15392

Very collegial group of scientists. Very good and flexible seminar organization

15421 – Rack-scale Computing | Dagstuhl Seminar | http://www.dagstuhl.de/15421

I absolutely loved the seminar. Great people, great scientists, great organization, great infrastructure, great staff. Initially I am thought that the location is not good but now I see that being remote is an advantage to make this seminar great – it brings the attendants literally together.

15441 – Duality in Computer Science | Dagstuhl Seminar | http://www.dagstuhl.de/15441

Best: Full-time interaction with participants, excellent facilities. Worst: nothing.

15021 – Concurrent computing in the many-core era | Dagstuhl Seminar | http://www.dagstuhl.de/15021

Dagstuhl = perfect

15041 – Model-driven Algorithms and Architectures for Self-Aware Computing Systems | Dagstuhl Seminar | http://www.dagstuhl.de/15041

I wish I had accepted earlier invitations to previous Dagstuhl workshops – really a wonderful environment for forming collaborations.

15031 – Understanding Complexity in Multiobjective Optimization | Dagstuhl Seminar | http://www.dagstuhl.de/15031

I found Dagstuhl a wonderful place for cross-fertilization.

15071 – Formal Foundations for Networking | Dagstuhl Seminar | http://www.dagstuhl.de/15071

The price for accommodation is AMAZING. The atmosphere for collaboration is AMAZING. Thank you so much for establishing Dagstuhl!

15211 – Theory of Evolutionary Algorithms | Dagstuhl Seminar | http://www.dagstuhl.de/15211

We were this time for the first time using the child care facility offered by Dagstuhl for our 20 months old daughter and it was really great: very smooth to organize and the person who took care of our daughter was perfect. So thank you for providing this child care. Note that we would have been happy to pay more for the child care, in case the price is a critical issue.

15222 – Human-Centric Development of Software Tools | Dagstuhl Seminar | http://www.dagstuhl.de/15222

I am so pleased to see Dagstuhl having more human-centered seminars (e.g., this one, the live coding seminar from a couple of years ago, the assessment of CS1 seminar next year). Computing is changing and broadening, and this is an important range of areas for Dagstuhl to support.

15222 – Human-Centric Development of Software Tools | Dagstuhl Seminar | http://www.dagstuhl.de/15222

Thanks for a wonderful experience! I really liked the "Dagstuhl culture" (unlocked doors, manor system).

15302 – Digital Scholarship and Open Science in Psychology and the Behavioral Sciences | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/15302

I feel very privileged to have been able to attend. The group's credentials and experience far outweighed mine but I rarely felt excluded or marginalized. It was a hell of an education and I will carry it into my work. Many attendees remarked that "this is what I thought academia was going to be like".

15351 – Computational Mass Spectrometry | Dagstuhl Seminar | http://www.dagstuhl.de/15351

Thanks go out to the organizers for a job well done, and for the Dagstuhl staff, especially the efficient and friendly catering staff.

15382 – Modeling and Simulation of Sport Games, Sport Movements, and Adaptations to Training | Dagstuhl Seminar | http://www.dagstuhl.de/15382

Dagstuhl is very unique; it is always a highlight to be here.

Resonanz zur Bibliographiedatenbank dblp

5.5

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 Dasgtuhl, oder durch die sozialen Medien. Darüber hinaus haben wir im November 2015 auf der dblp-Webseite öffentlich dazu aufgerufen, uns Meinungen und Kritiken zu dblp zukommen zu lassen.

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. Additionally, in November 2015, we issued a public call for comments via the dblp web site.

Lynda Hardman (President Informatics Europe) and Domenico Laforenza (President ERCIM)

dblp | http://dblp.org

Informatics Europe and ERCIM would like to emphasize the importance of this non-commercial European based service, available to all world-wide. In addition to locating authors and papers on research topics, the resource allows experimentation with metrics that can contribute to evaluations, allowing statistics to be based on a well-curated collection. The community is able to contribute to the collection, thus increasing the collection's data quality. The DBLP database has been used in research evaluations of faculties and institutes, as it is more reliable than other bibliographic data bases. It has also been used for analysis of co-author and topic patterns. Long-term studies on the development of the field depend on the continuity of the data collection. The personnel are very response to suggestions and requests, applying corrections suggested by individual users in a timely fashion. In short, DBLP is an excellent and unique resource allowing research to be carried out in the informatics and computer science and related fields and in monitoring their development.

Jochen L. Leidner (Director of Research, Thomson Reuters, London, UK)

dblp | http://dblp.org

Computer science research in the 21st century relies on a few select, yet very powerful tools: (i) the Google search engine (Google/Google Scholar), (ii) the ACM Digital Library, (iii) the DBLP computer science bibliography, and (iv) the Cornell pre-print archive (arXiv.org) are the four most prevalent sites I and my research group visit on a day-to-day basis to stay abreast of the latest emerging technology, and to research the related work for our own scientific publications. We would like to thank the DBLP team and its funders for the precious resource they have created, and continue to make available for free.

Ruben Ortega (CTO, Allen Institute for Artificial Intelligence, USA)

dblp | http://dblp.org

The Allen Institute for Artificial Intelligence launched our Semantic Scholar search engine (http://www.semanticscholar.org) on November 2, 2015. DBLP was integral to this launch as we are providing a full-text search engine across scholarly articles in computer science. DBLP is used as (1) a seed to our web crawlers, (2) the baseline for our weekly metrics reports, and (3) the "Gold Standard"for the metadata we use. Our launch happened faster and better because of our ability to bootstrap from the DBLP data set. As we look to expand our search engine into new scientific domains, we will need to work harder to assemble bibliographies that have as high a quality as DBLP.

Jan Mendling (Vienna University of Economics and Business, Austria)
dblp | http://dblp.org

DBLP is an invaluable resource that connects the global computer science community. It is the unique standard for understanding the contributions of international scholars and updates of conference and journal. As hosted by GI, Dagstuhl and the University of Trier, it establishes the role of the German Computer Science community as a central hub in this thriving worldwide research discipline.

anonymous email feedback
dblp | http://dblp.org

Without your fantastic service, I would not be able to do my job as it currently exists. This is not hyperbole — your open bibliographic information enables us to examine and conduct due diligence on new technologies, discover green-field markets and industries, even sourcing new investment opportunities. Our Research Department thoroughly enjoys and makes regular use of the journals, conference notes and publications in your database. We all sincerely hope that your efforts continue and expand in the future, and wish you nothing but success.

Kurt Mehlhorn (Max Planck Institute for Informatics, Saarbrücken, Germany)
dblp | http://dblp.org

DBLP is an extremely useful tool. I use it almost daily for tracking down individual papers and/or getting an overview of the research activity of a colleague. No other tool comes close.

Hubert Garavel (Inria LIG, Montbonnot-Saint-Martin, France)
dblp | http://dblp.org

DBLP is one of the most useful sites for my professional work. It is much more selective than Google Scholar and does not enter into the questionable game of bibliometric indexes. DBLP is really the *„Almanach de Gotha“* of computer scientists.

Jesper Larsson Träff (TU Wien, Austria)
dblp | http://dblp.org

It is absolutely vital for scientific work to have an accurate, comprehensive, trustworthy, public, non-commercial (e.g. Google only would be a dangerous option), independent literature database, and computer science is lucky to have the DBLP initiative! Coverage is extensive, and the information seems accurate and correct - and in rare cases where it is not, corrections are done swiftly. DBLP is an absolutely indispensable tool. Search functionality has improved over the years, more would of course be helpful (e.g. abstracts or keywords), but most important is that the database as is stays and is maintained.

Otthein Herzog (Jacobs University and University of Bremen, Germany)
dblp | http://dblp.org

If DBLP would not exist it would have to be invented! Thank you for many years of trusted service which allowed me to keep track of all publications interesting to me!

6 Endliche Automaten und das Unendliche
Finite Automata and the Infinite

Endliche Automaten und das Unendliche – und Dagstuhl als Katalysator
Finite Automata and the Infinite – and Dagstuhl as Catalysator

Vortrag zum Festakt "25 Jahre Schloss Dagstuhl", Schloss Dagstuhl, 3. Juli 2015
Talk at the Ceremony "25 Jahre Schloss Dagstuhl", Schloss Dagstuhl, July 3rd, 2015

Wolfgang Thomas (RWTH Aachen, Germany)

Prolog

Wie kann ich wohl am besten ausdrücken, was mir heute am Herzen liegt: Dagstuhl und seinem Team meinen tiefen Dank sagen für zweieinhalb Jahrzehnte phantastischer Arbeit? Und auf welche Weise kann ich vermitteln, wie dieses Paradies der Informatik-Forschung wirklich funktioniert? Und wie kann ich daneben auch wenigstens ein wissenschaftliches Ergebnis vorstellen (denn es soll ja ein wissenschaftlicher Vortrag sein), notwendigerweise aus meinem Fachgebiet, der theoretischen Informatik, freilich ohne dass man in Beweise einsteigen muss? Ich habe mir überlegt, dass ich über ein recht frühes Dagstuhl-Seminar berichten sollte. Es fand 1992 statt, wird unter der Nummer 9202 geführt, und es war das zweite Seminar, bei dem ich als Organisator beteiligt war. Allen, die damals dabei waren, steht es noch heute lebendig vor Augen.

Das Dagstuhl-Seminar 9202

The Dagstuhl-Seminar 9202 Als ich – auf dem Weg nach Dagstuhl – im Zug von Frankfurt nach Türkismühle einen Platz suchte, stieß ich zufällig ich auf ein Abteil, in dem zwei Leute saßen, die, so stellte sich bald heraus, Informatiker waren, und – mehr noch – ebenfalls auf der Reise nach Dagstuhl: Andrei Muchnik und Alexei Semenov, beide aus Moskau (siehe Abbildung 6.1).

Diese beiden Namen spielten auf der Einladungsliste für das Seminar eine ganz besondere Rolle – es war die Zeit kurz nach Öffnung der Grenzen zwischen dem Osten und dem Westen Europas, und es gab viel nachzuholen an wissenschaftlichem Austausch. Das Seminar in Dagstuhl sollte dazu einen Beitrag leisten. Wir werden sehen: Das gelang ganz hervorragend, und entscheidend dafür war Andrei Muchnik.

In Abbildung 6.2a sieht man aus dem ersten Band der Dagstuhl-Gästebücher die ersten Einträge der Seminarteilnehmer, bis hin zur Eintragung von Andrei Muchnik.

Andrei Muchnik war damals ein junger Wissenschaftler von 33 Jahren, Alexei Semenov sein Betreuer und Mentor. Semenov hatte auf Konferenzen auf die zahlreichen Resultate hingewiesen, die Muchnik und andere in Moskau erzielt hatten. Nun wollten wir in Dagstuhl Details erfahren.

Muchnik und Semenov hielten einen Vortrag gemeinsam. Die Zusammenfassung wurde von Semenov geschrieben. Abbildung 6.2b zeigt den Beginn aus dem Dagstuhl-Vortragsbuch.

Von den vielen Ergebnissen, die dann aufgelistet waren, wurde nur ein einziges im Vortrag näher ausgeführt, und zwar durch Muchnik. Es ist heute unter dem Namen "Muchniks Theorem" bekannt und stellte einen Durchbruch in der Methodik der (Programm-)Verifikation dar.

Prologue

How can I best express what I want to say today: to thank Dagstuhl and its team for two and a half decades of phantastic work? And how can I convey how this paradise of research in computer science really works? And how can I also present at least one scientific result (since I am supposed to give a scientific talk), necessarily from my own area of research, theoretical computer science, of course without entering proofs? I have decided to report on a rather early Dagstuhl-Seminar. It took place in 1992, is listed in Dagstuhl under the number 9202, and it was the second seminar in which I took part as one of the organizers. All who were present at the time still have vivid memories of this event.

The Dagstuhl-Seminar 9202

When I travelled to Dagstuhl, I looked for a place in the train from Frankfurt to Türkismühle and by chance took a seat next to two people who were, as it became clear very soon, computer scientists and – even more – also were heading towards Dagstuhl: Andrei Muchnik and Alexei Semenov, both from Moscow (see Figure 6.1).

These two names played a very special role on the list of invited researchers for the seminar – it was the time when the frontiers between the East and the West of Europe had been opened, and much was to be done to intensify scientific exchange. The seminar in Dagstuhl was supposed to contribute to this. As we shall see, this succeeded excellently, and essential for this success was Andrei Muchnik.

In Figure 6.2a the seminar page of the Dagstuhl Guest Book with the first names up to the entry of Andrei Muchnik is presented.

At the time Andrei Muchnik was a young scientist of 33 years, and Alexei Semenov was his supervisor and mentor. In conference talks, Semenov had pointed to many results of Muchnik and others from Moscow. Now in Dagstuhl we wanted to hear about details.

Muchnik und Semenov gave a talk together. The abstract was written by Semenov. Figure 6.2b shows the beginning of this text from the Dagstuhl Book of Abstracts.

From the numerous results that were listed there, only one, due to Andrei Muchnik, was presented in the talk. It is known today as "Muchnik's Theorem" and was a breakthrough in the methodology of (program-) verification.

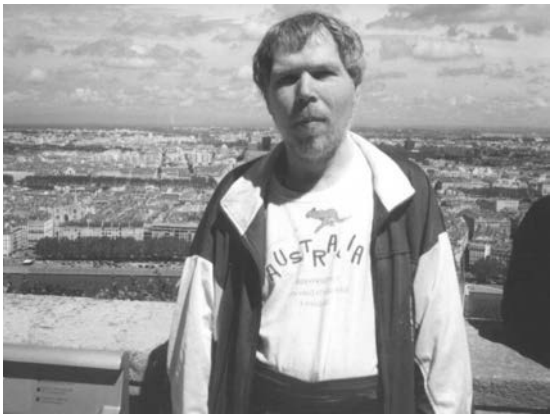

(a) Andrei A. Muchnik

(b) Alexei L. Semenov

Fig. 6.1
Andrei A. Muchnik und Alexei L. Semenov waren Teilnehmer beim Dagstuhl-Seminar "Automata Theory: Infinite Computations" (http://www.dagstuhl.de/9202), welches von Kevin Compton, Jean-Eric Pin und Wolfgang Thomas organisiert wurde und in der Woche vom 6.–10. Januar 1992 auf Schloss Dagstuhl stattgefunden hat.
Andrei A. Muchnik und Alexei L. Semenov were participants of the Dagstuhl-Seminar "Automata Theory: Infinite Computations" (http://www.dagstuhl.de/9202), organized by Kevin Compton, Jean-Eric Pin and Wolfgang Thomas and held in the week January 6–10, 1992, at Schloss Dagstuhl.

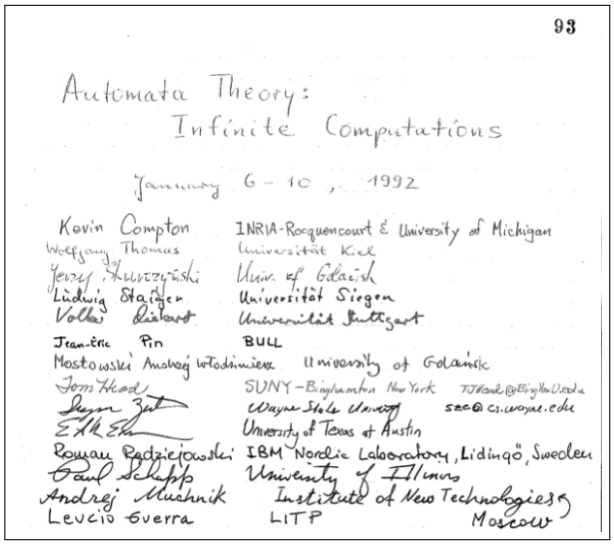

(a) Auszug aus dem Dagstuhl-Gästebuch

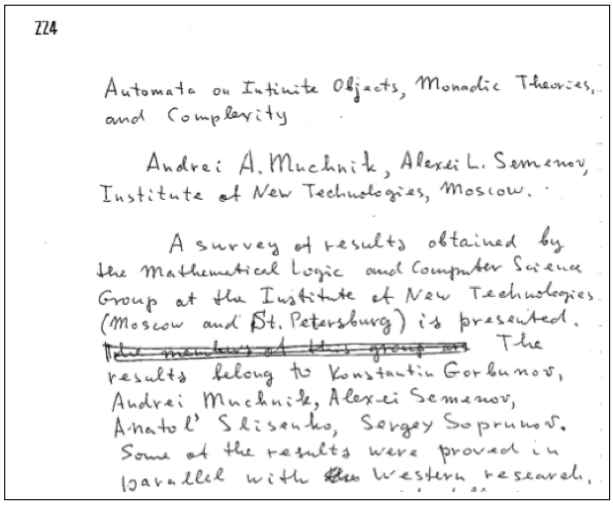

(b) Der handschriftliche Abstract von Alexei L. Semenov

Fig. 6.2
Auszüge aus den Dagstuhl-Büchern zum Aufenthalt von Andrei A. Muchnik und Alexei L. Semenov in Dagstuhl.
Excerpts form the Dagstuhl books on the stay of Andrei A. Muchnik und Alexei L. Semenov in Dagstuhl.

Das Resultat

In der heutigen Terminologie gehört der Satz von Muchnik zum Forschungsbereich des "Infinite-State Model Checking". Hier möchte man algorithmisch die Korrektheit von Programmen mit unendlichem Zustandsraum verifizieren. Im allgemeinen ist das natürlich unmöglich; das Halteproblem für Turing-Maschinen ist ja unentscheidbar (wir betrachten da die unendlich vielen Konfigurationen als "Systemzustände"). Aber in interessanten Fällen gelingt die automatische Verifikation. Der Satz von Muchnik ist ein Schlüssel dafür.

Allgemein haben wir es mit einem unendlichen Graphen G zu tun, dessen Knoten die Systemzustände und dessen Kanten die Übergänge von Zustand zu Zustand

The Result

In current terminology, Muchnik's Theorem belongs to the research area "Infinite-State Model Checking". The aim here is to verify in an algorithmic manner the correctness of programs with an infinite state-space. Of course, this is impossible in general since the halting problem for Turing machines is undecidable (we regard the infinitely many Turing machine configurations as "system states"). However, in interesting cases automatic verification is possible. Here Muchnik's Theorem is a key.

In general, we are dealing with an infinite graph G whose nodes represent the system states and whose edges the transitions from state to state. The correctness claim (or "specification") is expressed by a logical formula φ, and

darstellen. Die Korrektheitsbehauptung (oder "Spezifikation") wird durch eine Logik-Formel φ erfasst, und das Korrektheitsproblem ist die Frage, ob G die Formel φ erfüllt. In den Anwendungen, etwa der Protokollverifikation, zielt die Korrektheitsfrage allerdings nicht auf den Graphen selber, sondern auf alle möglichen Pfade durch den Graphen; das sind alle möglichen "Systemläufe". Gerade in der Protokollverifikation ist es vernünftig, diese Pfade als nicht-terminierend, also unendlich anzusehen (eine Termination bedeutet ja "Absturz des Systems"). Und normalerweise zeichnet man auch einen Zustand als Initialzustand i aus. Dann kann man alle Pfade durch G, die in i beginnen, als unendlichen Baum darstellen, dessen Wurzel der Knoten i ist. Jeder endliche Teil-Lauf von i aus führt zu einem eindeutig bestimmten Knoten dieses Baumes, und ein unendlicher Systemlauf entspricht einem unendlichen Pfad durch den Baum.

Ist ein Graph G mit Initialknoten i gegeben, so nennt man den entstehenden Baum $T(G, i)$ auch die "Abwicklung von G von i aus". Das ist eine fundamentale Begriffsbildung, die jedem in der Informatik geläufig ist. Hier ein kleines Beispiel für einen endlichen Graphen mit den Knoten 1, 2, 3 und Initialknoten 2. (Die Kanten, also die Zustandsübergänge, sind hier durch "Aktionennamen" a und b beschriftet.)

the correctness problem is the question whether G satisfies φ. In applications, for example in protocol verification, the correctness question is not about the graph itself but rather about the collection of all possible paths through the graph; these are the "system runs". In a context such as protocol verification it is appropriate to consider these paths as non-terminating and thus infinite (since termination means "system crash" here). And it is reasonable to designate one of the states as initial (denoted i). Then all paths through G that start in i can be presented as paths of an infinite tree with root i. Each finite partial run from i ends in a uniquely determined node of the tree, and an infinite system run determines an infinite path through the tree.

Given the graph G with initial node i, one calls the resulting tree $T(G, i)$ the "unfolding of G from i". This is a fundamental concept with which each computer scientist is familiar. A small example may serve as an illustration, where G is the finite graph with nodes 1, 2, 3 and initial node 2. (The edges, i.e., the transitions from state to state, are labelled with "action names" a and b.)

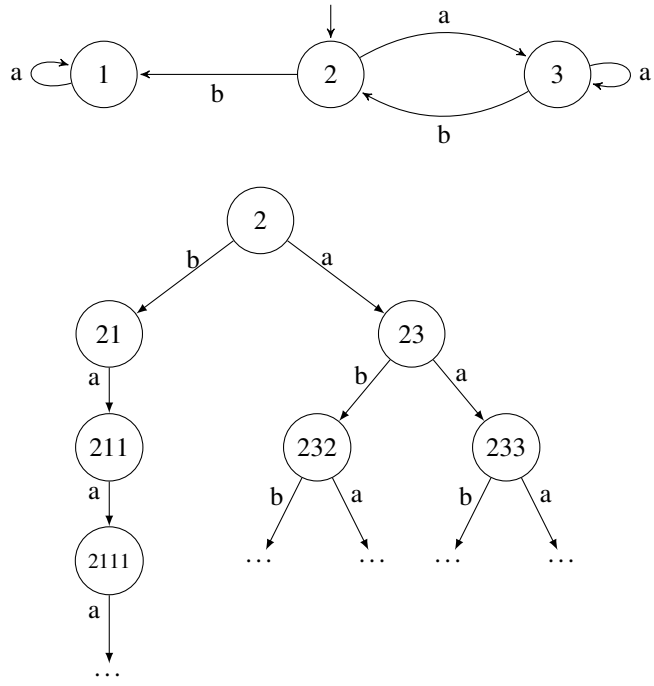

Die Verifikationsfrage lautet nun:

Gegeben ein Graph G mit Initialknoten i und eine Logikformel φ, gilt φ im Abwicklungsbaum $T(G, i)$?

Eine algorithmische Lösung hängt natürlich davon ab, wie kompliziert G ist und welche Logik man für die Formulierung von φ nimmt. Gerade bei der Logik muss man sehr vorsichtig sein, um nicht in den Morast der Unentscheidbarkeit zu geraten. Andererseits braucht man für die Anwendungen eine gewisse Ausdrucksstärke.

The verification problem can now be phrased as follows:

Given a graph G with initial node i and a logic formula φ, does φ hold in the unfolding $T(G, i)$?

Of course, an algorithmic solution depends on how complex G is and which logic is taken for the formulation of φ. Especially regarding the choice of the logic one has to be very careful in order to avoid ending up in the morass of undecidability. On the other hand, one needs some expressive power to cope with interesting applications.

Ein Logiksystem, das in diesem Sinne viele Vorteile bietet, ist die "MSO-Logik", offiziell "monadic second-order logic". In ihr hat man, wenn es um Graphen oder Bäume geht, Variablen $x, y, \ldots$ für Knoten und Variablen $X, Y, \ldots$ für Knotenmengen zur Verfügung, man kann z.B. durch $E_a(x, y)$ ausdrücken, dass von x nach y eine mit a beschriftete Kante führt, und man kann Formeln bilden mit den üblichen Booleschen Junktoren und mit den Quantoren $\exists$ und $\forall$, die über Knoten oder über Knotenmengen rangieren.

Zwei Beispiele über dem Bereich der Graphen mögen die Ausdruckmöglichkeiten dieser Logik illustrieren. Die Aussage "G hat einen Zykel der Mindestlänge 2" drückt man so aus (informell geschrieben):

> Es gibt verschiedene Knoten x und y mit Kante von x nach y, so dass man von y durch einen Pfad wieder x erreicht,

wobei die Existenz eines Pfades von y nach x so ausgedrückt wird:

> Jede Knotenmenge, die y enthält und die unter den "Kantenübergängen" abgeschlossen ist (jeweils von z mit $E(z, z')$ nach z'), muss auch x enthalten.

Ein anderes Beispiel betrifft die Existenz von Färbungen, z.B. die Eigenschaft "G ist 3-färbbar". Hier sagt man in MSO-Logik unmittelbar das, was die Definition ausmacht:

> Es gibt drei Mengen X,Y,Z, die eine Partition der Menge aller Knoten definieren und für die gilt, dass je zwei Knoten x,y, die durch eine Kante verbunden sind, zu zwei verschiedenen dieser Mengen gehören.

Nun fassen wir unser Verifikationsproblem bezogen auf die MSO-Logik wie folgt: Für welche Abwicklungsbäume T kann man die Frage

> Erfüllt T die MSO-Formel φ?

algorithmisch entscheiden? Ist dies für einen Baum T möglich, sagt man auch, die MSO-Theorie von T sei entscheidbar.

Betrachten wir ein ganz elementares Beispiel, den unendlichen binären Baum T_2 mit den beiden Kantenrelationen linker Nachfolger und rechter Nachfolger (vgl. die nachfolgende Figur). Man weiß: Die MSO-Theorie von T_2 ist entscheidbar. So einfach dieses Ergebnis klingt, es ist eines der schwierigsten Resultate der mathematischen Logik. Es wurde von Michael Rabin gegen Ende der 1960er Jahre bewiesen und ist ein Leuchtturmsatz, der seitdem in Tausenden von Arbeiten im Bereich der Informatik benutzt wurde. Seine Aufarbeitung in eine Form, die man in Vorlesungen präsentieren kann, dauerte Jahrzehnte (und bei diesem Projekt habe ich auch selbst mitgewirkt). Schlüsselidee für den Beweis ist die erstaunliche Tatsache, dass man Formeln in endliche Automaten übersetzen kann, die den Baum von der Wurzel aus längs aller Pfade parallel durchlaufen. Dies ist der entscheidende Schritt, um das "unendliche Problem" der Wahrheitswertbestimmung für eine Formel auf etwas Endliches zu reduzieren, was dann eine algorithmische Lösung erlaubt.

Der Satz von Muchnik ist eine weitreichende Verallgemeinerung des Satzes von Rabin. An Stelle eines einzigen Baumes, des binären Baumes T_2, tritt eine riesige Klasse von Bäumen, nämlich die Abwicklungen von Graphen mit entscheidbarer MSO-Theorie:

A logical system that combines many advantages in this respect is "MSO-logic", officially "monadic second-order logic". In this logic, referring to graphs or trees, one has variables $x, y, \ldots$ for nodes and variables $X, Y, \ldots$ for sets of nodes, one can express, for example, by $E_a(x, y)$ that there is an edge from x nach y labelled a, and one can build formulas with the usual Boolean junctors and the quantifiers $\exists$ und $\forall$ that range over nodes, respectively over sets of nodes.

Two examples over graphs may illustrate the expressive power of this logic. The statement "G has a cycle of length at least 2" is expressed as follows (written here informally):

> There are distinct nodes x and y with an edge from x to y such that there is a path from y back to x

where the existence of a path from y to x is expressed as follows:

> Each set of nodes that contains y and is closed under "edge transitions" (i.e., steps from z with $E(z, z')$ to z') also contains x.

Another example is about the existence of colorings, for example the property "G is 3-colorable". Here it is possible to say in MSO-logic just what the definition of 3-colorability gives:

> There are three sets X,Y,Z defining a partition of the set of all nodes such that any two nodes x, y that are connected by an edge belong to two different sets of X, Y, Z.

Now we can state our verification problem referring to MSO-logic: For which trees T obtained by unfolding can the question

> Does T satisfy the MSO-formula φ?

be decided algorithmically for any given φ? When this is possible for a tree T, one also says that the "MSO-theory of T is decidable".

Let us consider a very elementary example, the infinite binary tree T_2 with the two edge relations "left successor" and "right successor" (cf. the subsequent figure). It is known that the MSO-theory of T_2 is decidable. This statement may sound simple, but in fact it is one of the most difficult results in mathematical logic. It was proved by Michael Rabin around the end of the 1960's and is a landmark theorem that has since been used in thousands of papers in the area of computer science. The work-out in a form that can be presented in lectures took decades (and in this project I have also participated myself). The key idea for the proof is the astonishing fact that one can translate formulas into certain finite automata that traverse the tree under consideration from the root in a parallel fashion along all infinite paths. This is the crucial step to reduce the "infinite problem" of determining the truth value of a formula to something finite which then allows for an algorithmic solution.

Muchnik's Theorem is a far-reaching generalization of Rabin's result. Instead of a single tree, the binary tree T_2, a huge class of trees is covered, namely the unfoldings of all graphs with a decidable MSO-theory:

> If the MSO-theory of G is decidable and the initial node i definable in MSO-logic, then also the MSO-theory of the unfolding $T(G, i)$ is decidable.

Let us look at an example.

Ist die MSO-Theorie des Graphen G entscheidbar und der Initialknoten i in MSO-Logik definierbar, dann ist auch die MSO-Theorie des Abwicklungsbaums $T(G, i)$ entscheidbar.
Schauen wir uns wieder ein Beispiel an.

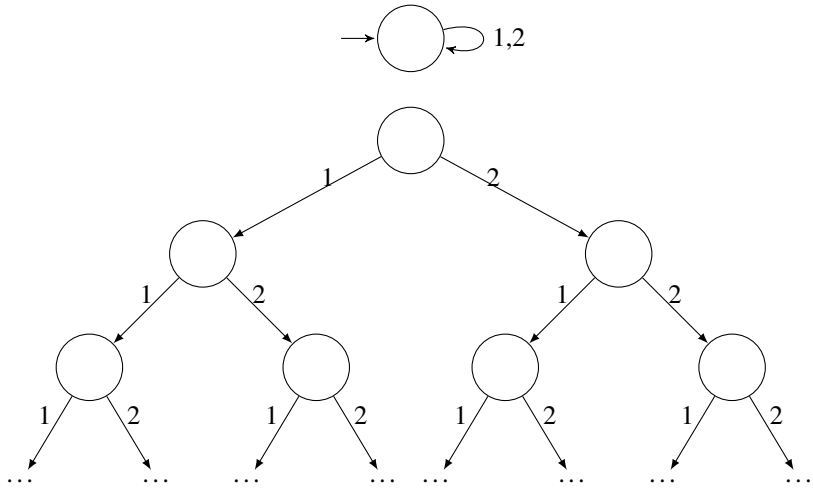

Als G nehmen wir den oben präsentierten Mini-Graphen, der aus einem einzigen Knoten i besteht, von dem eine 1- und eine 2-Kante zurück nach i führen. Natürlich ist die MSO-Theorie dieses Mini-Graphen entscheidbar – das ist so für jede Struktur, die nur aus einem einzigen Element (oder auch nur aus endlich vielen Elementen) besteht. Die Abwicklung dieses Minigraphen aber ist der unendliche binäre Baum! Der höchst komplexe Satz von Rabin ergibt sich also als winziger Spezialfall aus dem Satz von Muchnik.

Wir ersparen uns hier die schreckliche Arbeit, dieses Resultat im Einzelnen zu verstehen. Muchnik betrachtete sogar eine etwas komplexere Baumstruktur als die von uns benutzte Abwicklung eines Graphen (nämlich "tree iterations" an Stelle von "unfoldings"). Und wiederum war der Schlüssel für den Entscheidbarkeitsbeweis die Reduktion von MSO-Formeln auf geeignet konzipierte endliche Automaten, die unendliche Bäume von der Wurzel aus parallel durchlaufen. Dieser Zugang, das Unendliche mit endlichen Automaten in den Griff zu bekommen, ist – wie im Titel unseres Vortrags angedeutet – wieder der methodische Kern.

As G we take the above presented mini-graph consisting of a single node i, from which a 1-labelled and a 2-labelled edge point back to i. Of course the MSO-theory of this mini-graph is decidable – this holds for any structure with just one element (or with just finitely many elements). But the unfolding of this mini-graph is the infinite binary tree! Thus we obtain Rabin's highly complex theorem as tiny special case from Muchnik's Theorem.

We shall not enter here the terrible work to understand this result in detail. Muchnik even considered a more complex type of tree structure obtained from a graph (called "tree iteration" rather than "unfolding"). Again the key to establish the decidability proof is the reduction of MSO-formulas to suitably designed finite automata that traverse infinite trees in parallel along all the infinite paths. The approach to get hold of the infinite using finite automata – as indicated in the title of this talk – is again the methodological core idea.

Der Seminarvortrag

Im Dagstuhl-Seminar von 1992 wurde der Beweis für den Satz von Muchnik erstmalig der internationalen Community vorgestellt.

Der Vortrag im Seminarraum "Kaiserslautern" war ein Ereignis ganz besonderer Prägung. Vorne an der Tafel stand Andrei Muchnik, etwas introvertiert, mit einer Aura stiller aber sicherer Souveränität, wie sie eben ein Genie an sich hat, und neben ihm stand sein weltgewandter und zugleich besorgter Betreuer, Alexei Semenov. Der gemeinsame Auftritt wies beiden wohldefinierte Rollen zu: Muchnik sprach Russisch, und Semenov übersetzte und schrieb an die Tafel.

The Seminar Talk

In the Dagstuhl-Seminar of 1992 the proof of Muchnik's Theorem was presented to the international community for the first time.

The talk in the seminar room "Kaiserslautern" was a rather special event. In front, at the blackboard, Muchnik was standing, somewhat introverted, conveying an aura of silent but firm sovereignty – as it comes with a genius, and next to him was his internationally experienced and caring mentor, Alexei Semenov. The joint performance gave each of them a well-defined role: Muchnik spoke in Russian, and Semenov translated and wrote on the blackboard.

Das begann so: Muchnik sprach einen Satz auf Russisch. Darauf Semenov: "Let T be a tree." Darauf Muchnik in perfektem Englisch: "No: Let T be a *binary* tree." Das gespannte Publikum war höchst amüsiert, und für eine Weile verschob sich die Aufmerksamkeit von der Frage "Was trägt Muchnik vor?"' zu der profaneren Frage "Wie geht dieses Duett weiter?".

Irgendwie schafften es die beiden, zu einem schlüssigen Zusammenwirken zu finden, mit immer größerem Anteil von Muchnik. Und in den Stunden und Tagen nach dem Vortrag wurde weiter diskutiert, bis schließlich am Ende der Seminarwoche die allgemeine Überzeugung erreicht war, dass der Satz von Muchnik stimmt.

Die weitere Entwicklung – ein GI-Dagstuhl-Seminar

Obwohl der Satz von Muchnik nach dem Dagstuhl-Seminar in der "Szene" bekannt und akzeptiert war, dauerte es einige Zeit, bis ein vollständiger Beweis publiziert war. Muchnik selber verfasste kein Papier dazu; er arbeitete an immer neuen Fragen, zunehmend mit Verbindungen zur Komplexitätstheorie, und er überließ die Ausarbeitung seines Satzes anderen. Im Jahre 2007 kam dann die bestürzende Nachricht, dass dieser geniale und hochproduktive Forscher plötzlich verstorben war.

Ein erster vollständiger Beweis des Satzes von Muchnik wurde 1998 durch Igor Walukiewicz verfasst. Eine zugänglichere Darstellung wurde dann wieder in Dagstuhl erarbeitet, im Rahmen eines "GI-Dagstuhl-Seminars", das von Erich Grädel, Thomas Wilke und mir selbst organisiert wurde und bei dem etwa 20 Nachwuchswissenschaftler/innen wichtige aktuelle Ergebnisse aus Automatentheorie und Logik erarbeiteten. Diese GI-Dagstuhl-Seminare (in Kooperation mit der GI, der Gesellschaft für Informatik, eingerichtet) sind ein weiteres Juwel im Dagstuhl-Programm. Sie sammeln eine Community aktiver junger Leute, sie führen zu einer synthetischen und sorgfältigen Erarbeitung aktueller Ergebnisse, und als Resultat erscheint zumeist ein Band, der dann eine höchst nützliche Quelle für das jeweilige Forschungsgebiet ist. Die Beiträge unseres GI-Dagstuhl-Seminars (mit dem Thema "Automata, Logics, and Infinite Games") wurde denn auch im Jahre 2001 als Band 2500 bei den Springer Lecture Notes of Computer Science veröffentlicht[38], und ein wichtiges Kapitel, verfasst von Dietmar Berwanger und Achim Blumensath, erklärte den Satz von Muchnik [1]. Wer sich heute über den Satz von Muchnik informieren will, schaut dort nach.

Eine Modellhierarchie

Der Satz von Muchnik stieß eine Tür auf: Der Weg war frei, um eine reichhaltige Landschaft von unendlichen System-Modellen zu erschließen, für die man Verifikation automatisieren kann. Zu Anfang der 2000er Jahre wurde die Tür zu einem weiten Tor, als Didier Caucal vorschlug, den Satz von Muchnik mit einer weiteren Modellkonstruktion zu kombinieren. Hier wird nicht aus einem

This started as follows: Muchnik spoke a sentence in Russian. Then Semenov: "Let T be a tree". Then Muchnik, in perfect English: "No, let T be a *binary* tree." The attentive audience was highly amused, and for a while the interest shifted from the question "What does Muchnik present?" to the more profane question "How will this duet continue?".

Somehow the two managed to get to a really cooperative mode, with larger and larger parts taken by Muchnik. And in the hours and days after the talk the discussions continued, until finally at the end of the seminar week the general conviction was reached that Muchnik's Theorem is correct.

The Development After – a GI-Dagstuhl-Seminar

Although – following the Dagstuhl-Seminar – Muchnik's Theorem was known and accepted in the community, it took some time until a complete proof was published. Muchnik himself did not write a paper on this; he worked on many problems, more and more with connections to complexity theory, and left the work-out of his result to others. In 2007 the sad news were received that this genius and highly productive researcher had suddenly died.

A first complete proof of Muchnik's Theorem was written in 1998 by Igor Walukiewicz. A more accessible exposition was then again produced in Dagstuhl, in the framework of a "GI-Dagstuhl-Seminar", organized by Erich Grädel, Thomas Wilke, and myself, where about 20 young researchers worked on important recent results in the area of automata theory and logic. The GI-Dagstuhl-Seminars (established in cooperation with GI, the Gesellschaft für Informatik) are another jewel of the Dagstuhl program. They bring together a community of active young people, they lead to a synthetic and careful rework of recent results, and as an outcome often a volume is published which then serves as a highly useful reference in the respective field. The contribution of our GI-Dagstuhl-Seminar appeared (under the title "Automata, Logics, and Infinite Games") in 2001 as volume 2500 of the Springer Lecture Notes of Computer Science[38]. An important chapter, written by Dietmar Berwanger and Achim Blumensath, explained Muchnik's Theorem in detail [1]. Anybody who wants to know about Muchnik's Theorem today will look it up in that chapter.

A Hierarchy of Models

Muchnik's Theorem opened a door: The way was free to explore a rich landscape of infinite models of systems for which verification can be automated. Early in the 2000s this door widened to a large gate when Didier Caucal suggested to combine Muchnik's Theorem with another type of model construction. Here one does not generate a tree from a graph but conversely a graph from a tree. The idea is the

[38] http://dx.doi.org/10.1007/3-540-36387-4

Graphen durch Abwicklung ein Baum erzeugt, sondern umgekehrt aus einem Baum ein Graph. Die Idee besteht darin, innerhalb eines gegebenen Baumes T gewisse Knoten auszuwählen (es können auch alle Knoten sein) und dann zu beschreiben, wie die Kanten zwischen diesen Knoten gezogen werden sollen; natürlich geht dann im allgemeinen die Baumstruktur verloren und es entsteht ein Graph G. Diese Beschreibung muss durch Formeln in der MSO-Logik ausdrückbar sein; man spricht dann von einer "MSO-Interpretation von G in T." Eine solche Beschreibung einer neuen Struktur G in einer gegebenen Struktur T ist eine sehr gängige Modellkonstruktion, und ein einfaches Resultat besagt, dass die MSO-Theorie von G entscheidbar ist, falls bereits die MSO-Theorie von T entscheidbar ist.

Caucals Idee war nun, die beiden Modellkonstruktionen der Abwicklung (aus einem Graph entsteht ein Baum) und der MSO-Interpretation (aus einem Baum entsteht ein Graph) abwechselnd anzuwenden. Betrachten wir (in der folgenden Figur) ein Beispiel, beginnend mit einem Graphen aus vier Knoten, seiner Abwicklung (vom linken oberen Knoten aus) und der Interpretation, die alle vorhandenen Knoten und Kanten übernimmt und neue (mit d, e und f beschriftete) Kanten hinzufügt; deren MSO-Beschreibung ist sehr einfach.

choose certain nodes in a given tree T (these may as well be all nodes of T) and then to describe how to draw edges between these nodes; of course, the tree structure is then dissolved and one obtains a graph G. This description of nodes and edges is required to be expressible in MSO-logic; one then speaks of an "MSO-interpretation of G in T". Such a description of a new structure G within a given structure T is a standard model construction, and a simple result says that the MSO-theory of the new structure G is decidable if already the MSO-theory of T is decidable.

Caucal's idea was to apply the two model constructions of unfolding (from a graph a tree is obtained) and MSO-interpretation (from a tree a graph is obtained) in alternation. Let us look at an example (see the figure below), starting with a graph of four nodes, its unfolding (from the top-left node), and the interpretation that takes over all the given tree nodes and all tree edges and adds new edges labelled with d, e, f; their MSO-description is quite easy.

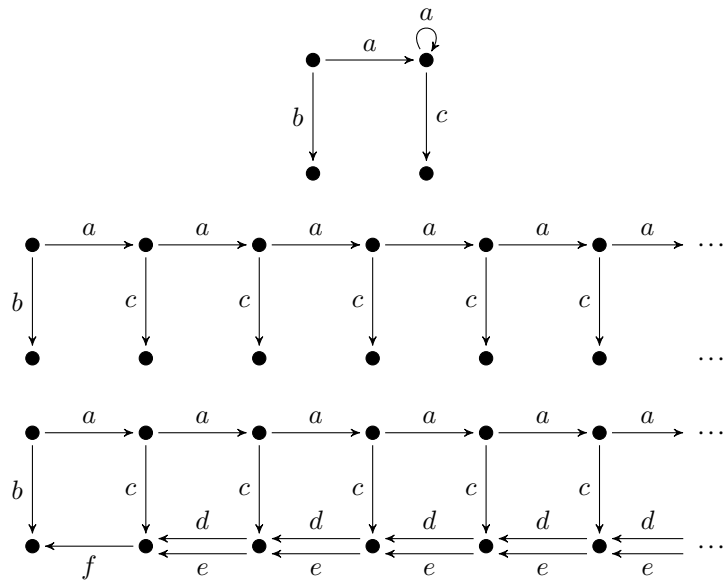

Durch immer neue Anwendung von Abwicklung und MSO-Interpretation erhält man die Strukturen der "Caucal-Hierarchie". Stufe 0 besteht aus den endlichen Bäumen und den endlichen Graphen, und Stufe $n+1$ entsteht aus Stufe n durch erneute Anwendung einer Abwicklung und einer MSO-Interpretation. Da jede endliche Struktur eine entscheidbare MSO-Theorie hat, ist nach dem Satz von Muchnik und der Bemerkung über MSO-Interpretationen die MSO-Theorie jeder Struktur der Caucal-Hierarchie entscheidbar. In unserem Beispiel gehört der endliche Graph zur Stufe 0, die beiden folgenden Strukturen gehören zur Stufe 1.

Schauen wir uns an, was durch eine weitere Abwicklung geschieht. Es entsteht eine interessante Baumstruktur der

By more and more applications of unfolding and MSO-interpretation one obtains the structures of the "Caucal hierarchy". Level 0 consists of the finite trees and the finite graphs, and level $n+1$ is obtained from level n by applying a further unfolding and an MSO-interpretation. Since each finite structure has a decidable MSO-theory, we know from Muchnik's Theorem and the above remark on MSO-interpretations that the MSO-theory of each structure in the Caucal hierarchy is decidable. In our example the finite graph belongs to level 0 and the other two structures to level 1.

Let us see what happens if we apply a further unfolding. We obtain an interesting tree structure of level 2 that is presented in the next figure (to be read without the dotted

Stufe 2, die in der folgenden Figur dargestellt ist (ohne die gestrichelte Linie zu lesen): Die mit d und e beschrifteten Doppelrückkanten des oben betrachteten Graphen fächern sich nun zu immer größeren Binärbäumen auf.

line): The double edges labelled d and e of the graph above split up into larger and larger finite binary trees.

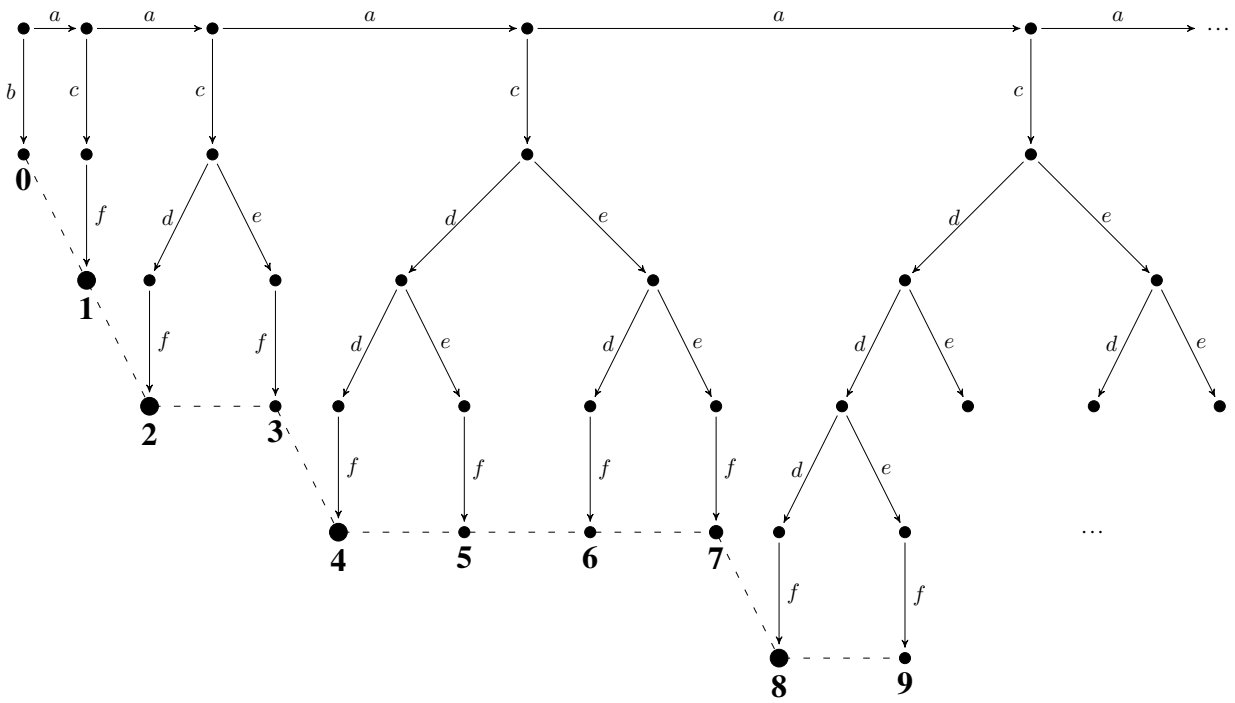

Nun beschreiben wir in diesem Baum durch eine MSO-Interpretation einen Graphen (der dann folglich auch zur Stufe 2 der Hierarchie gehört). Die Knotenmenge ist die Menge der Blätter des Baumes, und die Kanten werden vom einem Blatt zum jeweils nächsten Blatt gezogen (angedeutet durch die gestrichelte Linie). Wiederum gilt: Das kann man in der MSO-Logik leicht beschreiben. Wir erhalten eine Kopie der Nachfolgerstruktur $(N, Succ)$ der natürlichen Zahlen, in der Figur notiert mit $0, 1, 2, \ldots$, und mit der Nachfolgerrelation $Succ$, die jeweils von n nach $n+1$ führt. Das ist nun wirklich nichts Aufregendes. Aber wir können in der MSO-Logik auch eine Teilmenge der Blattknoten definieren, indem wir aus jedem Binärbaum das erste Blatt herausziehen. Diese Punkte sind fett markiert, und sie bilden, wie man sieht, die Menge der Zweierpotenzen. Es entsteht also die Struktur $(N, Succ, P)$ mit P = Menge der Zweierpotenzen. Die Struktur $(N, Succ, P)$ haben wir damit auf der Stufe 2 der Caucal-Hierarchie erhalten, und folglich ist ihre MSO-Theorie entscheidbar. Die Folge der Zweierpotenzen ist nicht mehr periodisch und nicht mehr etwas “Reguläres”, bei dem man die Entscheidbarkeit der MSO-Theorie einfach so erwarten kann. Dass man nun aber doch die Wahrheit von MSO-Aussagen in dieser Struktur algorithmisch entscheiden kann, ist bemerkenswert.

Wir haben damit einen ersten Schritt in eine Landschaft von Modellen getan, die wir beim gegenwärtigen Kenntnisstand noch gar nicht recht überblicken: Was sich jenseits von Stufe 2 so alles konstruieren lässt, ist ein Dickicht und noch nicht gut erforscht. Vor allem ist es die iterierte Anwendung des Satzes von Muchnik, die zu immer komplexeren Strukturen führt. Und für jede

Now we describe in this tree a graph by an MSO-interpretation (whence the graph also belongs to level 2 of the hierarchy). The set of graph nodes is the set of leaves of the tree, and the edges are taken to point from one leaf to the next (indicated by the dotted line). Again it is easy to describe this in MSO-logic. We obtain a copy of the successor structure $(N, Succ)$ of the natural numbers, denoted in the figure by $0, 1, 2, \ldots$, and with the successor relation $Succ$ that takes us from n to $n + 1$. Well, this is far from exciting. But we can describe in MSO-logic also a subset of the set of tree leaves, by extracting from each binary tree the first leaf. These leaves are marked boldface in the figure, and – as one observes – they form the set of powers of 2. So the structure $(N, Succ, P)$ with P = set of powers of 2 is generated. We have obtained $(N, Succ, P)$ on level 2 of the Caucal hierarchy, and hence the MSO-theory of this structure is decidable. The sequence of powers of 2 is not periodic and no more something “regular" where one can directly expect that the MSO-theory is decidable. That truth of MSO-statements in this structure can be decided algorithmically is remarkable.

With this example we have taken a first step into a landscape of models that – in the present state of knowledge – is not yet well understood. The class of models that can be constructed beyond level 2 is a jungle and not really explored. It is mainly the iterated application of Muchnik’s Theorem that leads to more and more complex structures. Each of these structures has a decidable MSO-theory whence “model checking" with respect to MSO-logic can be done algorithmically. Thus this hierarchy of models is a paradise for “infinite-state verification".

dieser Strukturen ist die MSO-Theorie entscheidbar, also das "Model Checking" bezüglich der MSO-Logik algorithmisch durchführbar. Damit ist diese Modellhierarchie ein Paradies für die "Infinite-State Verification".

Schlusswort

In vielen Dagstuhl-Seminaren wird heute das fortgesetzt, was mit dem Seminar vom Januar 1992 begann. Natürlich geht es nun oft darum, wie man aus der bloßen Tatsache, dass eine Theorie entscheidbar ist, praktisch nutzbare Algorithmen für die Verifikation ableiten kann. Hier hat es durch die Anstrengungen vieler Forschungsgruppen rund um den Erdball enorme Fortschritte gegeben, und die Vielfalt der Anwendungen entspricht dabei auch sehr gut der Vielfalt der Modelle, die wir durch Anwendung des Satzes von Muchnik erhalten.

Und wie bereits 1992 ist Dagstuhl bei diesen Anstrengungen auch weiterhin Katalysator für unsere Forschung.

In diesem Sinne sagen wir, die Forscherinnen und Forscher der Informatik:

Ad multas inventiones! Danke, Dagstuhl!

Conclusion

What started in the seminar of 1992 is continued these days in many further Dagstuhl-Seminars. Of course, a main objective is to convert the plain fact that a theory is decidable into practically useful algorithms for verification. The joint efforts of many research groups around the world have produced enormous progress, and the diversity of applications matches very well the diversity of models that are constructible via Muchnik's Theorem.

And as already in 1992, Dagstuhl continues to be in all these efforts a catalysator for our research.

And so we, the researchers of computer science, say:

Ad multas inventiones! Thank you, Dagstuhl!

Literaturverzeichnis

1 Dietmar Berwanger and Achim Blumensath, "The Monadic Theory of Tree-like Structures," in Erich Grädel, Wolfgang Thomas, and Thomas Wilke (eds.), *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, Lecture Notes in Computer Science, Vol. 2500, pp. 285–302, Springer, 2002. DOI: 10.1007/3-540-36387-4_16.

Picture credits

- Figure 6.1a: © Nikolay Vereshchagin. With kind permission of Nikolay Vereshchagin.
- Figure 6.1b: http://oxfordsymposium2014.gbsoxford.org.uk/images/speaker_images/Alexey.png
- Figures 6.2a and 6.2b: Archive Schloss Dagstuhl.

7 Die Seminare in 2015
The 2015 Seminars

Applications, Interdisciplinary Work

- Artifact Evaluation for Publications (15452)
- Computational Mass Spectrometry (15351)
- Computational Metabolomics (15492)
- Debating Technologies (15512)
- Design of Microfluidic Biochips: Connecting Algorithms and Foundations of Chip Design to Biochemistry and the Life Sciences (15352)
- Digital Scholarship and Open Science in Psychology and the Behavioral Sciences (15302)
- Dynamic Traffic Models in Transportation Science (15412)
- Evaluation in the Crowd: Crowdsourcing and Human-Centred Experiments (15481)
- Life-Long Health Behavior-Change Technologies (15262)
- Modeling and Simulation of Sport Games, Sport Movements, and Adaptations to Training (15382)
- Multi-Disciplinary Approaches to Reasoning with Imperfect Information and Knowledge – a Synthesis and a Roadmap of Challenges (15221)
- Self-Assembly and Self-Organization in Computer Science and Biology (15402)
- Smart Buildings and Smart Grids (15091)
- Social Concepts in Self-Organising Systems (15482)
- Software and Systems Traceability for Safety-Critical Projects (15162)

Artificial Intelligence, Computational Linguistics

- Artificial and Computational Intelligence in Games: Integration (15051)
- Formal Models of Graph Transformation in Natural Language Processing (15122)
- Multimodal Manipulation Under Uncertainty (15411)
- Normative Multi-Agent Systems (15131)
- Present and Future of Formal Argumentation (15362)
- The Mobile Revolution – Machine Intelligence for Autonomous Vehicles (15462)

Cryptography, Security, Privacy

- Genomic Privacy (15431)
- Quantum Cryptanalysis (15371)

Data Structures, Algorithms, Complexity

- Algorithms and Complexity for Continuous Problems (15391)
- Approximate and Probabilistic Computing: Design, Coding, Verification (15491)
- Complexity of Symbolic and Numerical Problems (15242)
- Computational Social Choice: Theory and Applications (15241)
- Empirical Evaluation for Graph Drawing (15052)
- Limitations of Convex Programming: Lower Bounds on Extended Formulations and Factorization Ranks (15082)
- Measuring the Complexity of Computational Content: Weihrauch Reducibility and Reverse Analysis (15392)
- Sparse Modelling and Multi-Exponential Analysis (15251)
- Symbolic Computation and Satisfiability Checking (15471)
- The Constraint Satisfaction Problem: Complexity and Approximability (15301)
- The Graph Isomorphism Problem (15511)
- Theory and Practice of SAT Solving (15171)
- Theory of Evolutionary Algorithms (15211)
- Understanding Complexity in Multiobjective Optimization (15031)

Databases, Information Retrieval, Machine Learning, Data Mining

- Approaches and Applications of Inductive Programming (15442)
- Cross-Lingual Cross-Media Content Linking: Annotations and Joint Representations (15201)
- Machine Learning with Interdependent and Non-Identically Distributed Data (15152)
- Mathematical and Computational Foundations of Learning Theory (15361)

Distributed Computation, Networks, Architecture, Systems

- Algorithms and Scheduling Techniques to Manage Resilience and Power Consumption in Distributed Systems (15281)
- Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations (15151)
- Concurrent Computing in the Many-Core Era (15021)
- Distributed Cloud Computing (15072)
- Formal Foundations for Networking (15071)
- Mixed Criticality on Multicore/Manycore Platforms (15121)
- Network Calculus (15112)
- Power-Bounded HPC Performance Optimization (15342)
- Quality of Experience: From Assessment to Application (15022)
- Rack-Scale Computing (15421)
- Secure Routing for Future Communication Networks (15102)

Geometry, Image Processing, Graphics, Visualization

- Bridging Information Visualization with Machine Learning (15101)
- Computational Geometry (15111)
- Holistic Scene Understanding (15081)
- The Message in the Shadow: Noise or Knowledge? (15192)
- Vision for Autonomous Vehicles and Probes (15461)

Software Technology, Programming Languages

- Advanced Stencil-Code Engineering (15161)
- Domain-Specific Languages (15062)
- Human-Centric Development of Software Tools (15222)
- Model-Driven Algorithms and Architectures for Self-Aware Computing Systems (15041)
- Programming with “Big Code” (15472)

Verification, Logic, Formal Methods, Semantics

- Challenges and Trends in Probabilistic Programming (15181)
- Circuits, Logic and Games (15401)
- Coalgebraic Semantics of Reflexive Economics (15042)
- Compositional Verification Methods for Next-Generation Concurrency (15191)
- Duality in Computer Science (15441)
- Information from Deduction: Models and Proofs (15381)
- Logics for Dependence and Independence (15261)
- Non-Zero-Sum-Games and Control (15061)
- Qualification of Formal Methods Tools (15182)
- Verification of Evolving Graph Structures (15451)

7.1 Concurrent Computing in the Many-Core Era

Organizers: Pascal Felber, J. Eliot B. Moss, Michael Philippsen, and Michael L. Scott

Seminar No. 15021

Date: January 4–9, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.1

Participants: José Nelson Amaral, Hagit Attiya, David F. Bacon, Annette Bieniusa, Hans-J. Boehm, Daniele Bonetta, Sebastian Burckhardt, Irina Calciu, Dave Dice, Stephan Diestelhorst, Sandhya Dwarkadas, Pascal Felber, Christof Fetzer, Maurice Herlihy, Antony Hosking, Milind Kulkarni, Viktor Leis, Charles E. Leiserson, Yossi Lev, Alexander Matveev, Maged M. Michael, J. Eliot B. Moss, Erez Petrank, Michael Philippsen, Torvald Riegel, Paolo Romano, Sven-Bodo Scholz, Michael L. Scott, Nir Shavit, Michael Swift, Gael Thomas, Osman Ünsal, Martin T. Vechev, Jons-Tobias Wamhoff

Context and Motivations

Thirty years of improvement in the computational power of CMOS uniprocessors came to an end around 2004, with the near-simultaneous approach of several limits in device technology (feature scaling, frequency, heat dissipation, pin count). The industry has responded with ubiquitous multi-core processors, but scalable concurrency remains elusive for many applications, and it now appears likely that the future will be not only massively parallel, but also massively heterogeneous.

Ten years into the multi-core era, much progress has been made. C and C++ are now explicitly parallel languages, with a rigorous memory model. Parallel programming libraries (OpenMP, TBB, Cilk++, CnC, GCD, TPL/PLINQ) have become mature enough for widespread commercial use. Graphics Processing Units support general-purpose data-parallel programming (in CUDA, OpenCL, and other languages) for a widening range of fields. Transactional memory appears likely to be incorporated into several programming languages. Software support is available in multiple compilers, and hardware support is being marketed by IBM and Intel, among others.

At the same time, core counts are currently lower than had once been predicted, in part because of a perceived lack of demand, and the prospects for increased core count over time appear to be constrained by the specter of dark silicon. Parallel programming remains difficult for most programmers, tool chains for concurrency remain immature and inconsistent, and pedagogical breakthroughs for the first- and second-year curriculum have yet to materialize. Perhaps most troublesome, it seems increasingly likely that future microprocessors will host scores or even hundreds of heterogeneous computational accelerators, both fixed and field-programmable. Programming for such complex chips is an exceptionally daunting prospect.

The goal of this Dagstuhl research seminar was to bring together leading international researchers from both academia and industry working on different aspects of concurrent computing (theory and practice, software and hardware, parallel programming languages, formal models, tools, etc.) in order to:

- assess the state of the art in concurrency, including formal models, languages, libraries, verification techniques, and tool chains;
- explore the many potential uses of emerging hardware support for transactional memory and synchronization extensions;
- envision next-generation hardware mechanisms;
- consider potential strategies to harness the anticipated explosion in heterogeneity; and
- investigate the interaction of synchronization and consistency with emerging support for low-latency byte-addressable persistent memory. (This last goal emerged late in the planning process, but became a major topic of discussion.)

Participants came from a wide variety of research communities, which seldom have the opportunity to meet together in one place. The seminar therefore provided a unique opportunity to focus diverse expertise on a common research agenda for concurrent computing on new generations of multi- and many-core systems.

Research Challenges

As part of this seminar, we specifically addressed the following challenges and open research questions, which are the focus of substantial investigation both in academia and in industry. These issues were addressed during the discussion at the workshop from the various perspectives of theory, concurrent algorithms, systems software, and microarchitecture.

The Future of Transactional Memory With the introduction this past year of TM-capable commodity processors

from IBM and Intel, TM research is increasingly turning to the question of how best to use the new hardware. What can and cannot be accomplished with the simple interfaces currently available? What might be accomplished with the addition of non-transactional loads and/or stores within transactions? (And how should such stores behave?) What support might be needed for nested transactions or nested parallelism?

Given that machines without TM will exist for many years, and that HTM will remain bounded by constraints on capacity, associativity, etc., how should hardware and software transactions interact? What hardware extensions might facilitate the construction of hybrid systems? Can hardware transactions be used to accelerate STM? Is TM hardware useful for purposes other than TM?

Beyond these basic questions, how do we integrate TM into the concurrency tool chain? How does one debug a black-box atomic operation? How should TM be embedded into programming languages? Should speculation be visible to the programmer, or should it be hidden within the implementation? How large can transactions reasonably become? Should they remain primarily a means of building concurrent data structures, or should they expand to encompass larger operations-even system-level functions like I/O, thread/process interactions, and crash recovery? As implementations proliferate, are there reasonable models of correctness that move beyond opacity? How should we benchmark TM code? What performance counters should future TM hardware provide to profilers? What kind of infrastructure is needed to perform regression testing of transactional code?

Heterogeneity GPUs are increasingly regarded as general-purpose computational resources, in platforms ranging from cell phones to supercomputers. Cell phones commonly include additional accelerators as well, for (de)compression, (de)encryption, and media transcoding. These and other accelerators (e.g., for linear algebra, pattern matching, XML parsing, or field-programmable functions) are likely to appear across the computing spectrum over the next few years.

In contrast to traditional (e.g., vector or floating-point) functional units, whose operations are uniformly short, and to traditional I/O devices, whose operations are uniformly long, accelerators can be expected to display a very wide range of response times. Long and variable response times suggest the need for resource management, to promote fair use across threads and applications. Short response times suggest the need for direct, user-level access-as already provided by GPU drivers from nVidia and (soon) AMD.

The prospect of contention for shared accelerators, accessed directly from user-level code, raises a host of questions for concurrent programming. How do we arbitrate shared access? Can traditional notions of locality be extended to accommodate heterogeneity? What happens to the tradeoff between local and remote computation when the alternatives use different instruction sets? What abstract models of progress/performance/time complexity are appropriate? Can operations that employ shared accelerators ever be considered non-blocking? How should we benchmark code that makes use of accelerators? What performance measures should heterogeneous architectures should provide to profilers? What kind of infrastructure is needed to perform regression testing in the face of heterogeneity?

Persistence Exceptions like magnetic core and battery-backed RAM notwithstanding, mainstream computing has long maintained a firm separation between fast, volatile working memory and slow, non-volatile (persistent) storage. Emerging low-latency, byte-addressable technologies like phase-change memory, memristors, and spin-torque-transfer memory bring this tradition into question. While near-term implementations may simply use low-latency nonvolatile memory as an accelerator for conventional file systems, alternative APIs may prove attractive. Specifically, it seems likely that future systems will give programmers the option of computing directly on persistent state, rather than reading it into working memory, using it there, and writing it out again. This possibility raises variants of many of the issues that have long concerned the concurrency community – consistency and atomicity in particular.

How should pointer-rich, non-file-based data be managed? Will we need automatic garbage collection? What will be the persistent analogues of nonblocking concurrent data structures? How will we ensure linearizability? Composability? A seemingly obvious option would add the 'D' (durability) to transactional memory's ACI (atomicity, consistency, and isolation). With little near-term prospect for integration of persistence and hardware TM, how will we minimize the overheads of persistent STM? What will the tradeoffs look like with respect to lock-based Lock-based programming models? What will be the division of labor between the operating system, runtime, and compiler? What will be the complexity models? Will we count "persistent accesses" the way we currently count remote memory accesses for concurrent objects in memory?

Pedagogy Once upon a time, concurrency was a specialized topic in the undergraduate curriculum, generally deferred to the operating systems course, or to an upper-level elective of its own. Now it is an essential part of the training of every computer scientist. Yet there is surprisingly little consensus on where it belongs in the curriculum, and how it ought to be taught. Alternatives range from "concurrency first," to infusion throughout the curriculum, to more extensive coverage in a more limited number of courses.

While the principal focus of the seminar was on research issues, participants had the opportunity to share both intuition and experience in the teaching of concurrency, during a dedicated panel session and as part of informal discussions. The following questions were notably discussed. What works, for which kinds of students? What languages and tool chains should we use? What textbooks do we need? What role (if any) should be played by deterministic parallel languages and constructs? Are there approaches, particularly for introductory students, that can offer parallel speedup for important applications, without the full complexity of the general case? Can these approaches reasonably be "staged" into intro-level courses?

Organization of the Seminar

The seminar lasted 5 days, each composed of short scientific presentations, with ample time for discussions, and break-out sessions during which various open questions were discussed in sub-groups. The first day of the seminar started with a general introduction and forward-looking presentations on concurrency and the challenges raised by heterogeneity and virtualization.

Ten technical sessions, with short presentations from the participants, took place during the seminar on:

- locks and TM;
- C++ status and standards;
- memory models;
- memory management and persistence;
- performance tuning and verification;
- distributed concurrency and fault-tolerance;
- thoughts on concurrency and parallelism;

- HW and portability;
- compilers, runtimes, and libraries; and
- languages and systems.

They were complemented by break-out sessions on “dealing with heterogeneity”, the “future of TM”, and “persistence”, as well as a plenary discussion on “virtualization”. Finally, a panel discussion was organized on the topic of “teaching concurrency”. The seminar concluded with an open discussion on the future of concurrency and the challenges that will need to be adddressed in coming years.

The topic of the sessions and their diversity illustrate the complexity of the challenges raised by concurrent computing on multi- and many-core systems. As one can expect from such prospective seminars, the discussions raised almost as many new questions as they provided answers on the addressed research challenges. Indeed, while there has been significant advances since the previous seminars (08241 and 12161), notably in terms of hardware support, few of the outstanding problems have been completely solved and new ones have emerged. For instance, hardware support for TM is now available in consumer CPUs but it cannot be used straightforwardly in real applications without relying on hybrid software/hardware strategies, notably to deal with the lack of progress guarantees and the possibility of spurious aborts.

As detailed in the rest of this report, the seminar has allowed the community to make significant progress on a number of important questions pertaining to concurrent computing, while at the same time defining a research agenda for the next few years. Participants provided very positive feedback following the seminar and expressed strong interest in follow-up events. Organizers strongly support the continuation of this series of seminars on concurrent computing, one of the most important and challenging fields in the era of multi- and many-core systems.

Fig. 7.1
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of Dr. Christian Lindig.

7.2 Quality of Experience: From Assessment to Application

Organizers: Katrien De Moor, Markus Fiedler, Peter Reichl, and Martín Varela

Seminar No. 15022

Date: January 4–7, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.57

Participants: Jan-Niklas Antons, Alemnew Asrese, Katrien De Moor, Philip Eardley, Sebastian Egger, Markus Fiedler, Farnaz Fotrousi, Pantelis Frangoudis, Samuel Fricker, Juan Pablo González Rivero, Poul Einar Heegaard, Tobias Hoßfeld, Kalevi Kilkki, Eirini Liotou, Toni Mäki, Sebastian Möller, Peter Reichl, Miguel Ríos Quintero, Werner Robitza, Henning Schulzrinne, Lea Skorin-Kapov, Samira Tavakoli, Christos Tsiaras, Astrid Undheim, Martín Varela, Min Xie, Patrick Zwickl

Within the past few years, Quality of Experience (QoE) has gone through an explosive growth and established itself as an independent, multidisciplinary field of research, both in academic and industrial communities. Significant advances have been made with respect to the conceptual understanding of QoE as well as in terms of methodology and instrumentation, and the earlier Dagstuhl seminars 09192 "From Quality of Service to Quality of Experience" and 12181 "QoE: From User Perception to Instrumental Metrics" have played a catalyzing role in this process: for example, by putting key challenges on the agenda, by stimulating (collaborative) activities that address them and by contributing to the establishment of a multi-disciplinary community around the topic, involving a range of actors with sometimes very different, yet complementary perspectives, priorities and motivations in relation to QoE. The main goal of this seminar was to strengthen and go beyond the current understanding on Quality of Experience (QoE) and its assessment, in order to address a logical yet highly challenging next step, namely to move from assessment to application and to translate insights in QoE and knowledge from this research field into forms of economic and/or societal "value". The main underlying motivation is that – even though the conceptual grounds and methodological implications of QoE are a very interesting and worthy research topic as such – they also represent milestones on the road to reach another ultimate goal: translating the theoretical and empirical understanding of QoE, its assessment and measures, into "value". This value can be rather explicit and concrete (e.g., increased revenue, or reduction of number of customer complaints), but it can also be intangible and more latent (e.g., customer loyalty, strengthened relation between a customer and a provider, enabling user empowerment, contributing to well-being).

The seminar brought together 27 participants to work towards this challenging goal. They were representing 13 different countries and 17 different institutions, resulting in a variety of different backgrounds and specific expertise domains. The seminar took place over 2.5 days and was organized in such a way that time for group discussion and interaction was maximized, while the time for individual presentations was kept to a minimum. At the beginning of the seminar, every participant was invited to write down three challenges fitting within this overall scope of the seminar. Thereupon, a concise presentation round was organized. Every participant was asked to make a short statement (5 minutes/1 slide) related to her or his challenges. These personal statements are included in the form of short abstracts in this report.

The main challenges and questions put forward by the participants were clustered on the fly into six broader topics, around which the seminar group work was organized, namely: "Theory and modeling", "QoE methodologies", "User factors and QoE", "QoE management", "Monetization of QoE" and "QoE in new domains". The group work was divided into two parts, with three topics being discussed in parallel in both parts of the group work. The initial assignment of participants to the six groups was deliberately organized randomly instead of thematically. The intention was to mix up participants with different backgrounds and interests as much as possible in order to stimulate open discussions and flow of thoughts. Participants had the possibility to switch to another group by exchanging with another participant in case they had a strong preference for another group. Every participant was involved in two discussion groups.

In between part 1 and 2 of the group work, a plenary reporting session was organized. During this plenary session, each group briefly presented the main points discussed and potential joint activities. During the final plenary reporting and closing session, the main points and outcomes from the second part of the group work were presented. Extensive summaries of the discussions and main outcomes for each of the six working groups are presented in full report. Due to the time constraints, there was unfortunately not enough time for deep follow-up discussions in the plenary

sessions. The seminar as such was also very briefly evaluated in the final plenary gathering. One important factor which would have further improved the participants' QoE and which was mentioned several times, is more time for "digestion" and "reflection" between the sessions (which was indeed limited, given the duration of the seminar). Overall however, and supported by the participants' feedback during and after the seminar, we can look back on a successful and productive seminar during which plans for several future and follow-up activities were made.

7.3 Understanding Complexity in Multiobjective Optimization

Organizers: Salvatore Greco, Kathrin Klamroth, Joshua D. Knowles, and Günter Rudolph

Seminar No. 15031

Date: January 11–16, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.96

Participants: Richard Allmendinger, Jürgen Branke, Dimo Brockhoff, Carlos A. Coello Coello, Salvatore Corrente, Matthias Ehrgott, Gabriele Eichfelder, Michael Emmerich, José Rui Figueira, Carlos M. Fonseca, Xavier Gandibleux, Martin Josef Geiger, Salvatore Greco, Jussi Hakanen, Carlos Henggeler Antunes, Hisao Ishibuchi, Johannes Jahn, Andrzej Jaszkiewicz, Yaochu Jin, Miłosz Kadziński, Kathrin Klamroth, Joshua D. Knowles, Renaud Lacour, Manuel López-Ibáñez, Luis Marté, Kaisa Miettinen, Sanaz Mostaghim, Vincent Mousseau, Mauro Munerato, Boris Naujoks, Luís Paquete, Silvia Poles, Robin Purshouse, Patrick M. Reed, Enrico Rigoni, Günter Rudolph, Stefan Ruzika, Serpil Sayın, Pradyumn Kumar Shukla, Roman Słowiński, Ralph E. Steuer, Michael Stiglmayr, Heike Trautmann, Tea Tusar, Daniel Vanderpooten, Simon Wessing, Margaret M. Wiecek, Xin Yao

Understanding complexity in multiobjective optimization is of central importance for the two communities, MCDM and EMO, and several related disciplines. It enables us to wield existing methodologies with greater knowledge, control and effect, and should, more importantly, provide the foundations and impetus for the development of new, principled methods, in this area.

We believe that a strong route to further progress in multiobjective optimization is a determination to understand more about the various ways that complexity manifests itself in multiobjective optimization. We observe that in several fields, ranging from engineering to medicine to economics to homeland security, real-world problems are very often characterized by a high degree of complexity deriving from the presence of many competitive objectives to be optimized, many stakeholders expressing conflicting interests and the presence of many technical parameters being unstable in time and for which we have imperfect knowledge. These very complex problems require a specific methodology, mainly based on multiobjective optimization, that, using high computational capacities, takes into account robustness concerns and allows an effective participation of the several stakeholders in the decision process.

The seminar took place January 11th–16th 2015. The main goals of the seminar were the exploration and elucidation of complexity in three fundamental domains:

Focus 1: Complexity in preference This topic is mainly concerned with elicitation, representation and exploitation of the preference of one or more users, for example: discovering and building preferences that are dynamic and unstable, group preference, complex structure of criteria,non-standard preferences, learning in multiobjective optimization.

Focus 2: Complexity in optimization This topic is mainly concerned with the generation of alternative candidate solutions, given some set of objective functions and feasible space. The following topics are examples for the wide range of issues in this context: high-dimensional problems, complex optimization problems, simulation-based optimization and expensive functions, uncertainty and robustness, interrelating decision and objective space information.

Focus 3: Complexity in applications An all-embracing goal is to achieve a better understanding of complexity in practical problems. Many fields in the Social Sciences, Economics, Engineering Sciences are relevant: E-government, Finance, Environmental Assessment, E-commerce, Public Policy Evaluation, Risk Management and Security issues are among the possible application areas.

During the seminar the program was updated on a daily basis to maintain flexibility in balancing time slots for talks, discussions, and working groups. The working groups were established on the first day in highly interactive fashion: at first each participant was requested to write her/his favorite topic on the black board, before a kind of collaborative clustering process was applied for forming the initial five working groups, some of them splitting into subgroups later. Participants were allowed to change working groups during the week, but the teams remained fairly stable throughout. Abstracts of the talks and extended abstracts of the working groups can be found in subsequent chapters of this report.

Further notable events during the week included: (i) a session devoted to discuss the results and the perspectives of this series of seminars after ten years of the first one, (ii) a hike within a time slot with worst weather conditions during the week, (iii) a presentation session allowing us to share details of upcoming

events in our research community, and (iv) a wine and cheese party made possible by a donation of UCL's *EPSRC Centre for Innovative Manufacturing in Emergent Macromolecular Therapies* represented by Richard Allmendinger.

Outcomes

The outcomes of each of the working groups can be seen in the sequel. Extended versions of their findings will be submitted to a Special Issue on "Understanding Complexity in Multiobjective Optimization" in the *Journal of Multi-Criteria Decision Analysis* guest-edited by the organizers of this Dagstuhl seminar.

This seminar resulted in a very insightful, productive and enjoyable week. It has already led to first new results and formed new cooperation, research teams and topics. In general, the relations between the EMO and MCDM community were further strengthened after this seminar and we can expect that thanks to the seminar a greater and greater interaction will be developed in the next few years.

Acknowledgements. Many thanks to the Dagstuhl office and its helpful and patient staff; huge thanks to the organizers of the previous seminars in the series for setting us up for success; and thanks to all the participants, who worked hard and were amiable company all week. In the appendix, we also give special thanks to Salvatore Greco as he steps down from the organizer role.

7.4 Model-Driven Algorithms and Architectures for Self-Aware Computing Systems

Organizers: Samuel Kounev, Xiaoyun Zhu, Jeffrey O. Kephart, and Marta Kwiatkowska

Seminar No. 15041

Date: January 18–23, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.164

Participants: Tarek F. Abdelzaher, Artur Andrzejak, Christian Becker, Steffen Becker, Kirstie Bellman, Nelly Bencomo, Sara Bouchenak, Javier Camara, Giuliano Casale, Lydia Y. Chen, Ada Diaconescu, Lukas Esterle, Antonio Filieri, Francesco Gallo, Kurt Geihs, Holger Giese, Sebastian Götz, Lars Grunske, Henry Hoffmann, Paola Inverardi, Alexandru Iosup, K. R. Jayaram, Evangelia Kalyvianaki, Joost-Pieter Katoen, Jeffrey O. Kephart, Samuel Kounev, Anne Koziolek, Heiko Koziolek, Marta Kwiatkowska, Philippe Lalanda, Chris Landauer, Peter Lewis, Martina Maggio, Ole Mengshoel, Arif Merchant, Aleksandar Milenkoski, Anders Robertsson, Felix Maximilian Roth, Hartmut Schmeck, Evgenia Smirni, Katinka Wolter, Xin Yao, Xiaoyun Zhu, Andrea Zisman

Seminar Description

Self-aware computing systems are best understood as a subclass of autonomic computing systems. The term, autonomic computing, was first introduced by IBM in 2001. Expressing a concern that the ever-growing size and complexity of IT systems would soon become too difficult for human administrators to manage, IBM proposed a biologically-inspired solution. An analogy was drawn between the autonomic nervous system, which continually adjusts the heart and respiratory rates, pupil dilation, and other lower-level biological functions in response to conscious decisions made by individuals, and autonomic computing systems, which are envisioned as managing themselves in accordance with high-level objectives from humans.

In an effort to enlist the academic community in a worldwide effort to meet this grand challenge, Kephart and Chess laid out a vision of autonomic computing in an IEEE Computing article in 2003 [1]. The article postulated a multi-agent architecture for autonomic computing systems consisting of interacting software agents (called autonomic elements) that consume computational resources and deliver services to humans and to other autonomic elements, and used that architecture as a structure against which a diverse set of research challenges were defined. One of the major challenges from a scientific perspective was the definition of appropriate abstractions and models for understanding, controlling, and designing emergent behavior in autonomic systems. Many different components of IT systems could be autonomic elements – database management systems, load balancers, provisioning systems, anomaly detection system, etc. In addition to managing their own behavior in accordance with policies established by humans or other autonomic elements, they also manage their relationships with other autonomic elements.

The self-managing properties of autonomic computing systems, including self-optimization, self-configuration, self-healing and self-protection, are expected to arise not just from the intrinsic self-managing capabilities of the individual elements, but even more so from the interactions among those elements, in a manner akin to the social intelligence of ant colonies. Understanding the mapping from local behavior to global behavior, as well as the inverse relationship, was identified as a key condition for controlling and designing autonomic systems. One proposed approach was the coupling of advanced search and optimization techniques with parameterized models of the local-to-global relationship and the likely set of environmental influences to which the system will be subjected.

In the ensuing decade, there has been much research activity in the field of autonomic computing. At least 8000 papers have been written on the topic, and explicit solicitations for papers on autonomic computing can be found in the call for papers of at least 200 conferences and workshops annually, including the International Conference on Autonomic Computing (ICAC), now in its tenth year. The European government has funded autonomic computing research projects for several million euros via the FP6 and FP7 programs, and the US government has funded research in this field as well.

In a retrospective keynote at ICAC 2011, Kephart assessed the state of the field, finding through bibliometric analysis that progress in the field has been good but uneven [2]. While there has been a strong emphasis on self-optimization in its many forms, there have been considerably fewer works on other key autonomic properties such as self-configuration, self-protection and self-healing. An apparent reason for this imbalance is that benchmarks that quantify these properties and allow them to be compared across different systems and methods are still largely lacking. Another finding was that much work remains to be done at the system level. In particular, while there has been considerable success in using machine learning and feedback control techniques to create adaptive autonomic elements, few authors have successfully built autonomic computing systems

containing a variety of interacting adaptive elements. Several authors have observed that interactions among multiple machine learners or feedback loops can produce interesting unanticipated and sometimes destructive emergent behaviors; such phenomena are well known in the multi-agent systems realm as well, but insufficiently understood from a theoretical and practical perspective.

It is worth noting that there is a substantial sub-community within the autonomic computing field that applies feedback control to computing systems. FeBID (Feedback Control Implementation and Design in Computing Systems and Networks), a key workshop in this space, began in 2006 as a forum for describing advances in the application of control theory to computing systems and networks. In 2012, FeBID acquired a new name (Feedback Computing) to reflect a much broader and colloquial interpretation of "feedback", in which the goals are no longer merely set points, and system models are not merely used to help transform or transduce signals, but may themselves be adapted through learning. The evolution of this sub-community of autonomic computing reflects a growing acceptance of the idea that, for an autonomic computing element or system to manage itself competently, it needs to exploit (and often learn) models of how actions it might take would affect its own state and the state of the part of the world with which it interacts.

Self-Aware Computing Systems

To understand how self-aware computing systems fit within the broader context of autonomic and feedback computing, we started with the following definition [4, 5] in the beginning of the seminar:

► **Definition.** A computing system is considered to be "self-aware" if it possesses, and/or is able to acquire at runtime, the following three properties, ideally to an increasing degree the longer the system is in operation:

- Self-reflective: Aware of its software architecture, execution environment, and hardware infrastructure on which it is running as well as of its operational goals (e.g., quality-of-service requirements, cost- and energy-efficiency targets),
- Self-predictive: Able to predict the effect of dynamic changes (e.g., changing service workloads) as well as predict the effect of possible adaptation actions (e.g., changing system configuration, adding/removing resources),
- Self-adaptive: Proactively adapting as the environment evolves in order to ensure that its operational goals are continuously met.

The three properties in the above definition are obviously not binary, and different systems may satisfy them to a different degree, however, in order to speak of "self-awareness", all three properties must apply to the considered system.

To realize the vision of "self-aware" computing systems, as defined above, we advocated a holistic model-based approach where systems are designed from the ground up with built-in self-reflective and self-predictive capabilities, encapsulated in the form of online system architecture models. The latter are assumed to capture the relevant influences (with respect to the system's operational goals) of the system's software architecture, its configuration, its usage profile, and its execution environment (e.g., physical hardware, virtualization, and middleware). The online models are also assumed to explicitly capture the system's operational goals and policies (e.g., quality-of-service requirements, service level agreements, efficiency targets) as well as the system's adaptation space, adaptation strategies and processes.

Figure 7.2 presents our vision of a self-aware system adaptation loop based on the MAPE-K control loop [3] in combination with the online system architecture models used to guide the system adaptation at runtime. In the following, we briefly describe the four phases of the adaptation loop.

Phase 1 (Observe/Reflect): In this phase, the managed system is observed and monitoring data is collected and used to extract, refine, calibrate, and continuously update the online system models, reflecting the relevant influences that need to be captured in order to realize the self-predictive property with respect to the system's operational goals. In the context of this phase, expertise from software engineering, systems modeling and analysis, as well as machine learning, is required for the automatic extraction, refinement and calibration of the online models based on observations of the system at runtime.

Phase 2 (Detect/Predict): In this phase, the monitoring data and online models are used to analyze the current state of the system in order to detect or predict problems such as SLA violations, inefficient resource usage, system failures, network attacks, and so on. Workload forecasting combined with performance prediction and anomaly detection techniques can be used to predict the impact of changes in the environment (e.g., varying system workloads) and anticipate problems before they have actually occurred. In the context of this phase, expertise from systems modeling, simulation, and analysis, as well as autonomic computing and artificial intelligence, is required to detect and predict problems at different time scales during operation.

Phase 3 (Plan/Decide): In this phase, the online system models are used to find an adequate solution to a detected or predicted problem by adapting the system at runtime. Two steps are executed iteratively in this phase: i) generation of an adaptation plan, and ii) prediction of the adaptation effects. In the first step, a candidate adaptation plan is generated based on the online models that capture the system adaptation strategies, taking into account the urgency of the problem that needs to be resolved. In the second step, the effects of the considered possible adaptation plan are predicted, again by means of the online system architecture models. The two steps are repeated until an adequate adaptation plan is found that would successfully resolve the detected or predicted problem. In the context of this phase, expertise from systems modeling, simulation, and analysis, as well as autonomic computing, artificial intelligence, and data center resource management, is required to implement predictable adaptation processes.

Phase 4 (Act/Adapt): In this phase, the selected adaptation plan is applied on the real system at runtime. The actuators provided by the system are used to execute the individual adaptation actions captured in the adaptation plan. In the context of this phase, expertise from data center resource management (virtualization, cluster, grid and cloud computing), distributed systems, and autonomic computing, is required to execute adaptation processes in an efficient and timely manner.

Broader Notion of Self-aware Computing

As a result of the working group "Defining Self-aware Computing Systems", a broader notion of self-aware computing was formulated:

► **Definition.** Self-aware computing systems are computing systems that:

1. *learn models* capturing *knowledge* about themselves and their environment (such as their structure, design, state, possible actions, and run-time behavior) on an ongoing basis and
2. *reason* using the models (for example predict, analyze,

consider, plan) enabling them to *act* based on their knowledge and reasoning (for example explore, explain, report, suggest, self-adapt, or impact their environment)

in accordance with *higher-level goals*, which may also be subject to change.

For a detailed discussion of the interpretation of this definition, we refer the reader to Section 4.1 of the full report.

Industrial Relevance

The envisioned novel algorithms and architectures for self-aware computing systems are of high relevance to the real-world problems faced by software developers and practitioners in the IT industry. Even though many of the specific problems have been researched upon within the aforementioned disciplines and communities, we believed the timing is right for adopting a broader integrated and interdisciplinary approach and exploiting synergies in the existing modeling and management approaches. The demand and the urgency for providing practical model-driven solutions to the described problems have never been higher, for the following reasons:

Large-scale, on-demand infrastructure: Although the cloud computing concept has been around for a long time, it wasn't until the last few years did we see a wide availability and adoption of cloud computing platforms. Such platforms provide infrastructure-on-demand to business critical applications and high performance computing workloads. Such highly dynamic, demand-driven environments make many existing automation schemes in computing systems inadequate, because they are mostly rule-based or heuristics-driven and cannot self-adapt to changes in both the infrastructure and the workloads.

Applications and workloads: The ever-increasing variety and complexity of modern applications and their workloads are placing more stress on computing systems and making many traditional management approaches obsolete. This is exacerbated by the extensive use of mobile devices and applications by an increasing population that produces new usage patterns and resource requirements.

Sensors and data: The numbers and types of sensors deployed in computing systems have never been greater, which lead to an explosion of runtime monitoring data that accurately capture the operating conditions of systems and software. Such data significantly enhance the chances for computing systems to Observe/Reflect (Phase 1) and to extract/refine/calibrate online system models that were difficult to learn otherwise, making a model-driven approach more feasible and reliable.

Need for automation: The IT industry is crying out ever so loud for automation technologies to help deal with the above challenges. Automation also helps reduce manual labor cost in management and administration and addresses the increasing gap between the number of skilled IT professionals and the industrial demand. There have been a growing number of startup companies that aim at developing automation solutions for capacity planning, provisioning and deployment, service level assurance, anomaly detection, failure/performance diagnosis, high availability, disaster recovery, and security enforcement. More research on modern-driven algorithms and architectures for self-aware computing can really feed into this new wave of innovations.

Organization of the Seminar

As inspired by the above described vision and approach towards its realization, we believed that the design of self-aware computing systems calls for an integrated interdisciplinary approach building on results from multiple areas of computer science and engineering including: i) software and systems engineering; ii) systems modeling, simulation and analysis; iii) autonomic and organic computing, machine learning and artificial intelligence; iv) data center resource management including virtualization, cluster, grid and cloud computing. This was the motivation of the research seminar. The list of invitees was carefully composed to provide a balance among these fields including both theoretical and applied research with participation from both academia and industry. We note that, in reality, each of the four mentioned communities is in fact comprised of multiple separate sub-communities although they have some overlap in their membership. While they can be seen as separate research communities, we consider them related in terms of their goals, with the difference being mostly in the specific focus of each sub-community and the employed scientific methods. The final participants of the seminar included representatives from each sub-community such that we cover the different relevant focus areas and scientific methodologies.

Achievements of the Seminar

This seminar has achieved its original goal of bringing together scientists, researchers, and practitioners from four different communities, including Software Engineering, Modeling and Analysis, Autonomic Computing, and Resource Management, in a balanced manner. The seminar program provided a basis for exchange of ideas and experiences from these different communities, offered a forum for deliberation and collaboration, and helped identify the technical challenges and open questions around self-aware computing systems. In summary, its achievements are mainly in the following two areas.

Identification of Synergies and Research Questions

By bringing together researchers from the above research fields and their respective communities, we avoid duplication of effort and exploit synergies between related research efforts.

During the seminar, we identified the following research questions and challenges that are of common interest to multiple communities:

- Design of abstractions for modeling quality-of-service (QoS) relevant aspects of systems and services deployed in dynamic virtualized environments. The abstractions should make it possible to capture information at different levels of detail and granularity allowing to explicitly model the individual layers of the system architecture and execution environment, context dependencies, and dynamic system parameters.
- Automatic model extraction, maintenance, refinement, and calibration during operation. Models should be tightly coupled with the system components they represent while at the same time they should abstract information in a platform-neutral manner.
- Efficient resolution of context dependencies including dependencies between the service deployment context and input parameters passed upon invocation, on the one hand, and resource demands, invoked third-party services, and control flow of underlying software components, on the other hand.
- Automatic generation of predictive models on-the-fly for online QoS prediction. The models should be tailored to

answering specific online QoS queries. The model type, level of abstraction and granularity, as well as the model solution technique, should be determined based on: i) the type of the query (e.g., metrics that must be predicted, size of the relevant parts of the system), ii) the required accuracy of the results, iii) the time constraints, iv) the amount of information available about the system components and services involved.
- Efficient heuristics exploiting the online QoS prediction techniques for dynamic system adaptation and utility-based optimization.
- Novel techniques for self-aware QoS management guaranteeing service-level agreements (SLAs) while maximizing resource efficiency or minimizing energy cost.
- Standard metrics and benchmarking methodologies for quantifying the QoS- and efficiency-related aspects (e.g., platform elasticity) of systems running on virtualized infrastructures.

The above research questions and challenges were considered in the context of our holistic model-based approach and the self-aware system adaptation loop presented in the previous section. Answering these questions can help determine what system aspects should be modeled, how they should be modeled, how model instances should be constructed and maintained at runtime, and how they should be leveraged for online QoS prediction and proactive self-adaptation.

The online system models play a central role in implementing the four phases of the described system adaptation loop. The term "'model'" in this context is understood in a broad sense since models can be used to capture a range of different system aspects and modeling techniques of different type and nature can be employed (e.g., an analytical queuing model for online performance prediction, a machine learning model for managing resource allocations, a statistical regression model capturing the relationship between two different system parameters, a descriptive model defining an adaptation policy applied under certain conditions). At the seminar, we advocate a model-based approach that does not prescribe specific types of models to be employed and instead we use the term "online system models" to refer to all information and knowledge about the system available for use at runtime as part of the system adaptation loop. This includes both descriptive and predictive models.

Descriptive models describe a certain aspect of the system such as the system's operational goals and policies (quality-of-service requirements and resource efficiency targets), the system's software architecture and hardware infrastructure, or the system's adaptation space and adaptation processes. Such models may, for example, be described using the Meta-Object-Facility (MOF) standard for model-driven engineering, heavily used in the software engineering community.

Predictive models are typically applied in three different contexts: i) to predict dynamic changes in the environment, e.g., varying and evolving system workloads, ii) to predict the impact of such changes on system metrics of interest, iii) to predict the impact of possible adaptation actions at runtime, e.g., application deployment and configuration changes. A range of different predictive modeling techniques have been developed in the systems modeling, simulation and analysis community, which can be used in the "detect/predict" phase of our adaptation loop, e.g., analytical or simulative stochastic performance models, workload forecasting models based on time-series analysis, reliability and availability models based on Markov chains, black-box models based on statistical regression techniques. Finally, models from the autonomic computing and machine learning communities can be used as a basis for implementing the "plan/decide" phase of our adaptation loop. Examples of such models are machine learning models based on reinforcement learning or analytical models based on control theory.

Two important goals of the seminar were to discuss the applicability of the various types of models mentioned above in the context of self-aware computing systems, and to evaluate the tradeoffs in the use of different modeling techniques and how these techniques can be effectively combined and tailored to the specific scenario. As discussed above, in each phase of the self-aware adaptation loop, multiple modeling techniques can be employed. Depending on the characteristics of the specific scenario, different techniques provide different tradeoffs between the modeling accuracy and overhead. Approaches to leverage these tradeoffs at runtime in order to provide increased flexibility will be discussed and analyzed.

Finally, the practical feasibility and associated costs of developing system architecture models was also extensively discussed. We also identified a major target of future research in the area of self-aware computing, which is to automate the construction of online system models and to defer as much as possible of the model building process to system runtime (e.g., the selection of a suitable model to use in a given online scenario, the derivation of adequate model structure by dynamically composing existing template models of the involved system components and layers, the parameterization of the model, and finally, the iterative validation and calibration of the model). Such an approach has the potential not only to help reduce the costs of building system architecture models, but also to bring models closer to the real systems and applications by composing and calibrating them at runtime based on monitoring of the real observed system behavior in the target production environment when executing real-life operational workloads.

Impact on the Research Community

By bringing together the aforementioned four communities, the research seminar allowed for cross-fertilization between research in the respective area. It has raised the awareness of the relevant research efforts in the respective research communities as well as existing synergies that can be exploited to advance the state-of-the-art of the field of self-aware computing systems. The seminar has left to this Dagstuhl Report that provides an up-to-date point of reference to the related work, currently active researchers, as well as open research challenges in this new field. Given that a significant proportion of the proposed participants are from industry, the seminar also fostered the transfer of knowledge and experiences in the respective areas between industry and academia.

In addition to producing this joint report summarizing, we also found enough support and interest among the seminar participants to continue the collaboration through the following venues: i) writing a joint book to publish at Springer with chapter contributions from the seminar participants, ii) establish a new annual workshop on self-aware computing to provide a forum for exchanging ideas and experiences in the areas targeted by the seminar.

Overall, the seminar opened up new and exciting research opportunities in each of the related research areas contributing to the emergence of a new research area at their intersection.

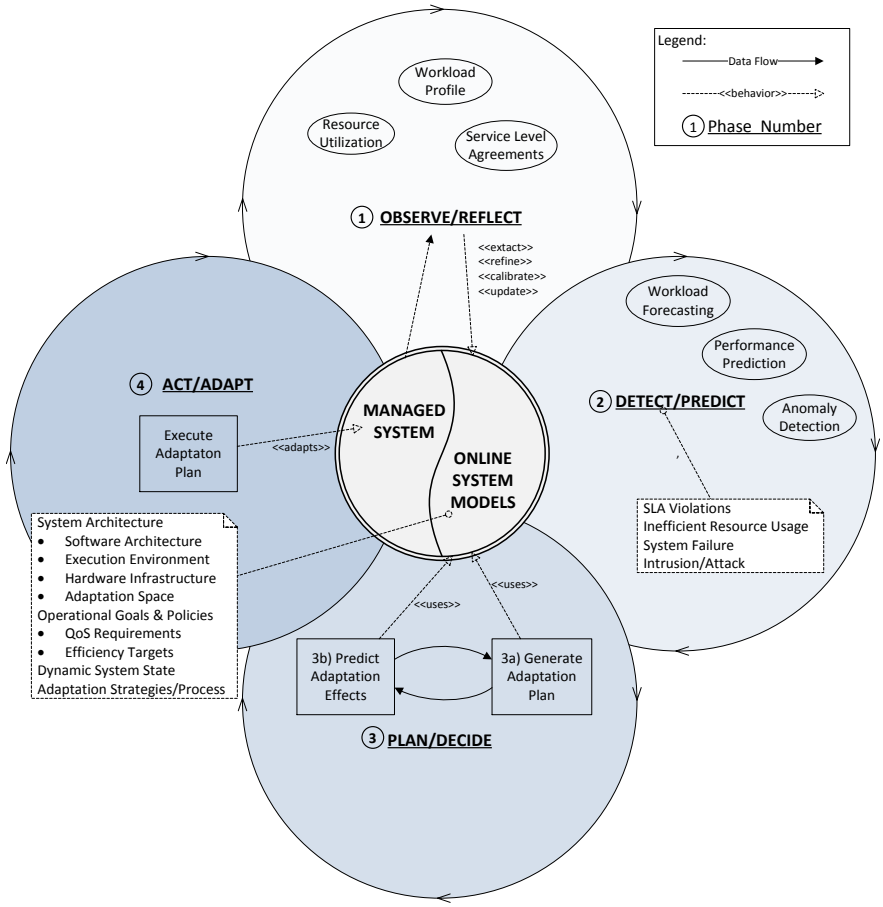

Fig. 7.2
Self-aware system adaptation loop.

References

1 Jeffrey O. Kephart, David M. Chess, "*The Vision of Autonomic Computing*," in IEEE Computer, 36(1):41–50, 2003. DOI: 10.1109/MC.2003.1160055

2 Jeffrey O. Kephart, "*Autonomic Computing: The First Decade*", in Proc. of the 8th ACM Int'l Conf. on Autonomic Computing (ICAC'11), pp. 1-2, ACM, 2011. DOI: 10.1145/1998582.1998584

3 IBM Corporation, "*An Architectural Blueprint for Autonomic Computing*", IBM White Paper, 4th Edition, 2006.

4 Samuel Kounev, "*Engineering of Self-Aware IT Systems and Services: State-of-the-Art and Research Challenges*", in Proc. of the 8th European Performance Engineering Workshop (EPEW'11), LNCS, Vol. 6977, pp. 10–13, Springer, 2011. DOI: 10.1007/978-3-642-24749-1_2

5 Samuel Kounev, Fabian Brosig, and Nikolaus Huber, "*Self-Aware QoS Management in Virtualized Infrastructures*", in Proc. of the 8th ACM Int'l Conf. on Autonomic Computing (ICAC'11), pp. 175–176, ACM, 2011. DOI: 10.1145/1998582.1998615

7

7.5 Coalgebraic Semantics of Reflexive Economics

Organizers: Samson Abramsky, Alexander Kurz, Pierre Lescanne, and Viktor Winschel

Seminar No. 15042

Date: January 18–21, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.197

Participants: Samson Abramsky, Achim Blumensath, Filippo Bonchi, Neil Ghani, Helle Hvid Hansen, Michael Hauhs, Julian Hedges, Alexander Kurz, Stéphane Le Roux, Pierre Lescanne, Fabio Mogavero, Paulo Oliva, Prakash Panangaden, Daniela Petrisan, Marcus Pivato, Jan Rutten, Martin Scheffel, Heiner Schumacher, Alexandra Silva, Baltasar Trancon y Widemann, Elias Tsakas, Evguenia Winschel, Viktor Winschel, Philipp Zahn

A growing number of researchers have been discovering analogies in the foundations of both computer science and economics. The goal of this seminar is to interface computer science with economics and game theory and to take advantage of the programming language semantics methods in theoretical computer science based on lambda calculus, coalgebras, modal logic and category theory.

The theoretical thread of interest to this seminar and common to both computer science and economics is the phenomenon that may be circumscribed by notions such as reflexivity, self-reference, impredicativity, infinite regress, recursion, or fixed points.

In computer science, the phenomena of self-reference, self--application and recursion played a crucial role in the foundational work of Gödel, Church, Turing and Kleene in the 1930s. Nevertheless, powerful mathematical models of the semantics of recursion became available only with the work of Scott on models of the untyped lambda calculus and subsequent research in domain theory. The combination of domain theory with the theory of types in programming languages and their categorical semantics has led to the development of a powerful tool box. More recently, this tool box has been further strengthened by advances in coalgebra. It provides for a wide variety of dynamic systems the mathematical tools of (bi)simulation and coinduction as well as a variety of techniques from category theory.

In economics, and the social sciences more generally, reflexivity arises from the obvious fact that cognitive agents reason about themselves, others and the society they live in. This leads to self-reference and recursion in, for example, theories of belief formation as beliefs of beliefs (Harsanyi type spaces) or theories of institutions as rules to change rules. More generally, a social system consists of individuals who are learning about a process in which others are learning as well. Learning the state of an interactive system is therefore rather different compared to learning the parameter values that govern a physical process. When the observer is a part of the system, the act of learning changes the thing to be learned. The traditional mathematical tools in economics are hardly suited to solve these problems in a sufficiently general way and they make it difficult for computer scientists, once they need to solve similar or common problems, to understand the problem formulation and the solutions already achieved by economists.

The specific subfields of computer science and economics discussed above suggest to explore methods of program semantics and category theory in general and, in particular, of bisimulation and coinduction in economics. Furthermore, coalgebra gained prominence as providing models for concurrency, a topic that has hardly been touched upon in economics explicitly, even so it underlies the most general kind of issues in economics, namely those regarding centralization versus decentralization in theories of economic systems, administration, firms and markets.

Particular topics in which we see scope for methods from the semantics of programming languages include infinitely repeated games, econometrics and system theory, epistemic game theory and interactive learning in multi-agent systems.

More generally, research in program semantics and logics in computer science is typically motivated by problems arising in programming languages and software engineering. In one direction, economic modeling will become more important in software engineering. In the other direction, computational economics may as well profit from a modern approach to language design not only in terms of reflexivity at the theoretical level but also at the practical level of modeling software.

7.6 Artificial and Computational Intelligence in Games: Integration

Organizers: Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius

Seminar No. 15051

Date: January 25–30, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.207

Participants: Dan Ashlock, Christian Bauckhage, Rafael Bidarra, Bruno Bouzy, Michael Buro, Alex J. Champandard, Simon Colton, Michael Cook, Peter I. Cowling, Mirjam P. Eladhari, Ian Horswill, Graham Kendall, Pier Luca Lanzi, John M. Levine, Antonios Liapis, Simon M. Lucas, Brian Magerko, Michael Mateas, Joshua Allen McCoy, Mark J. Nelson, Ana Paiva, Mike Preuss, Günter Rudolph, Spyridon Samothrakis, Tom Schaul, Noor Shaker, Moshe Sipper, Adam M. Smith, Gillian Smith, Pieter Spronck, Kenneth O. Stanley, Ruck Thawonmas, Tommy Thompson, Julian Togelius, Michael Treanor, Marc van Kreveld, Santiago Ontanon Villar, Mark Winands, Georgios N. Yannakakis, R. Michael Young, Fabio Zambetta, Jichen Zhu, Alex Zook

The research field of artificial and computational intelligence in games focuses on the wide variety of advanced computational techniques used to create, enhance, and improve the experiences of humans interacting with and within virtual environments. By its nature the field is broad and multidisciplinary. People working in it include academic researchers from a variety of disciplines, corporate researchers in the games industry as well as in other industries, and independent game developers. The methods used include symbolic AI techniques such as reasoning, constraint-solving and partial-order planning as well as biologically-inspired techniques such as evolutionary computation and neural networks, statistical techniques such as support vector machines and clustering, as well as special-purpose techniques such as behavior trees. These are applied to games ranging from board and card games to first-person shooters and real-time strategy games as well as abstract mathematical games, and are used to, for instance, play games, model players, tell stories, generate levels and other game content, and match players. Different researchers have different goals, including developing and testing new AI methods, creating interesting and believable non-player characters, improving the game production pipeline, studying game design through computational means, and understanding players and patterns of interaction. Often several goals overlap in the same project.

Recent years have seen considerable progress in several of the techniques used in the field, as well as rapid development in what kind of research questions are asked and what kind of games are studied with which methods. It has become increasingly clear that many of the research goals require a multidisciplinary approach, or at least a combination of techniques that, in the past, were considered separate research topics. For instance, with respect to the behavior of virtual agents, ten years ago researchers mainly aimed at making such behavior more "effective," which can often be achieved with straightforward computational reasoning. Nowadays, however, researchers aim at making the behavior of virtual agents more "natural" in their interaction with humans, which requires contributions not only from computer science, but also from psychology and social sciences, and which requires a wide variety of techniques, such as player modeling, adaptation, reasoning, and computational linguistics.

To move the research field forward, it is therefore of crucial importance to facilitate the integration of the disciplines and techniques that are involved in this research. The various strands, methodological approaches, and research directions need to inform each other and collaborate, to achieve a whole that is more than the sum of its parts. The goal of the second *Dagstuhl Seminar on Computational and Artificial Intelligence in Games* was to explicitly take the first steps along this path of integration, and investigate which topics and techniques would benefit most from collaboration, how collaboration could be shaped, and which new research questions may potentially be answered.

The seminar was held between January 25 and January 30, 2015. To stimulate interaction between the participants, which is essential in this context, the seminar was structured around workgroups rather than presentations. The organizers started the seminar on Monday morning with a series of brief presentations on potential discussion topics, after which the participants formed their own workgroups around a variety of topics, not necessarily those brought up by the organizers. Workgroups typically consisted of 3 to 10 people from different backgrounds, who worked together for no more than one day. At regular intervals workgroups reported on their findings in a plenary session, after which new workgroups were formed.

At the start of the seminar it was announced that Thursday would be set aside for practical work. Participants could use that day to implement some of the ideas that had come up in the previous days, in the form of a game, a competition, a design document, or a research proposal. While the organizers

deliberately gave the participants the option to simply continue with the workgroups if they so wished, all participants actually got involved in the practical work, some of them even working on multiple projects in parallel.

The results of the workgroups and the practical sessions are briefly related in the remainder of these proceedings. The 13 abstracts on workgroups cover automated and AI-based game design; game analytics; interdisciplinary research methods; design of believable characters; general video game playing; creativity facet orchestration; methods and formal design for procedural content generation; design of "fun" gameplaying bots; communication on game AI research, computers that play like humans; and neural networks for games. The 11 abstracts on practical sessions cover the Planet Wars competition; the automatic generation of games, mazes, and text; Twitter bots; sonification of character reasoning; MCTS and representation learning for procedural content generation; two AI-based games; and the design for a board game.

A special issue of the *IEEE Transactions on Computational and Artificial Intelligence in Games* will be published on the topic of this Dagstuhl Seminar. While this issue is open for submission for any researcher in this field, it is expected that several of the workgroups of the seminar will submit papers on their results.

As organizers, we knew that the first seminar that we organized in 2012 was considered a great success, and we had expected more people to accept our invitations for this second seminar than for the previous one. However, demand for attending the seminar was even greater than we expected. Almost everyone we first invited immediately accepted our invitation. Moreover, everybody who accepted their invitation indeed showed up at the seminar. We were forced by capacity concerns to not invite many people who, by their strength of contribution in the field, should have been present. We are certain that we could easily have doubled the number of participants visiting the seminar, and that each of those participants would have made a strong contribution.

The value of these Dagstuhl Seminars is indisputable. Considering the large number of researchers that should be invited to a seminar that attempts to cover the whole, very broad research field of Computational and Artificial Intelligence in Games, we believe that it is wise for a future seminar to narrow down the topic, so that it can be restricted to a smaller number of participants that are active in the selected subfield. Naturally, considering the fact that "integration" is such an important issue in the research field, care must be taken to ensure that every discipline interested in and involved in the subfield is represented.

7.7 Empirical Evaluation for Graph Drawing

Organizers: Ulrik Brandes, Irene Finocchi, Martin Nöllenburg, and Aaron Quigley

Seminar No. 15052

Date: January 25–30, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.1.243

Participants: Daniel Archambault, Ulrik Brandes, Markus Chimani, Emilio Di Giacomo, Walter Didimo, Tim Dwyer, Martin Fink, Irene Finocchi, Rudolf Fleischer, Seok-Hee Hong, Andreas Karrenbauer, Michael Kaufmann, Karsten Klein, Stephen G. Kobourov, Giuseppe Liotta, Catherine C. McGeoch, Tamara Mchedlidze, Wouter Meulemans, Tamara Munzner, Lev Nachmanson, Martin Nöllenburg, Yoshio Okamoto, Maurizio Patrignani, Huamin Qu, Aaron Quigley, Bernice E. Rogowitz, Ignaz Rutter, Falk Schreiber, Janet Siegmund, Bettina Speckmann, Matthias F. Stallmann, Michael Stoffel, Dorothea Wagner, Alexander Wolff, Kai Xu

Graph Drawing provides, among other things, the algorithmic foundations for network information visualization. It has considered implementation and experimentation as integral aspects from its very inception and recent research has demonstrated varying approaches to empirical evaluation. Experimental standards, however, have never been established, and little progress toward higher levels of sophistication can be observed.

The seminar was a community effort organized as a hands-on training event. It brought together experts on experimentation from fields with an established experimental tradition (referred to as "trainers"), and a group of graph drawing researchers expected to act as exponents and multipliers ("participants"). After two days of invited lectures on experimental methodology in different disciplines and a problem selection session, participants spent three days in working groups designing experiments. Trainers moving between groups and intermittent reporting session facilitated knowledge dissemination.

Participant feedback in the Dagstuhl survey indicates that the inclusion of trainers was highly appreciated. A number of experimental designs for a broad range of problems have been developed, and it is expected that many of them will be implemented and carried out in collaborative follow-up work.

As everyone who has ever been to Schloss Dagstuhl knows, Dagstuhl seminars are the ideal forum for achieving such goals. The fact that a considerable part of the graph drawing community came together for a week to focus on experimentation is expected to lead to a rapid diffusion of the seminar results and foster the acceptance of new methodology and criteria within the community.

On behalf of all participants, the organizer express their sincere gratitude to the Dagstuhl staff for their outstanding service and support.

Fig. 7.3
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of (above) Dr. Christian Lindig and (below) TreeState Productions GmbH, www.treestate.de.

7.8 Non-Zero-Sum-Games and Control

Organizers: Krishnendu Chatterjee, Stéphane Lafortune, Nicolas Markey, and Wolfgang Thomas

Seminar No. 15061

Date: February 1–6, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.1

Participants: Nathalie Bertrand, Dietmar Berwanger, Patricia Bouyer-Decitre, Romain Brenguier, Benedikt Brütsch, Véronique Bruyère, Krishnendu Chatterjee, Laurent Doyen, Rüdiger Ehlers, John Fearnley, Gilles Geeraerts, Hugo Gimbert, Alessandro Giua, Axel Haddad, Thomas Anton Henzinger, Rasmus Ibsen-Jensen, Barbara Jobstmann, Jan Krčál, Jan Křetínský, Stéphane Lafortune, Kim Guldstrand Larsen, Simon B. Laursen, Christof Löding, Nicolas Markey, Peter Bro Miltersen, Benjamin Monmege, Necmiye Ozay, Nicolas Perrin, Sophie Pinchinat, Mickael Randour, Jean-François Raskin, Dorsa Sadigh, Ocan Sankur, Sven Schewe, Anne-Kathrin Schmuck, Armando Solar-Lezama, Jiří Srba, John G. Thistle, Wolfgang Thomas, Ufuk Topcu, Stavros Tripakis, Martin Zimmermann

Games played on graphs provide the framework to study a wide range of problems that are central in computer science, for example, reactive synthesis, well-formedness of systems, checking compatibility of behavorial type systems, etc. The traditional study of games has been for two-player zero-sum perfect-information deterministic games with Boolean objectives (where a win of one player coincides with a loss by the other player). Fundamental results of this theory are contributions of automata theory that go back to the 1960's (Büchi, McNaughton, Rabin).

Significant progress has been achieved in the last few decades both in terms of theoretical results (understanding the complexity of such games, developing efficient algorithms) as well as their practical applicability (in reactive synthesis and controller synthesis). The current research directions explore several important extensions of the traditional study, namely, multi-player games, games with partial-observation, quantitative aspects in games, as well as application of game results in other domains. In this regard, the connection to control theory is important: The methodology of "supervisory control" has developed in parallel to the emergence of the game-theoretic approach, and a joint and integrating view of these closely related branches of research seemed overdue.

The Dagstuhl Seminar "Non-Zero Sum Games and Control" addressed these developments, with a particular emphasis on the connections to control theory. The response to the call for participation was very positive, and the 42 scientists joining the seminar represented the full range of topics mentioned above. There was a very good mixture between young and "established" researchers, and the participation of female researchers (making a quarter) was high (in the context of computer science). In order to support the understanding between the different research branches, it was decided to have on each half-day at least one survey talk, outlining a field and describing general challenges; some of these talks were contributed by young researchers. The speakers of the survey talks were Rüdiger Ehlers, Tom Henzinger, Barbara Jobstmann, Stéphane Lafortune, Kim Larsen, Peter Bro Miltersen, Jean-François Raskin, Armando Solar-Lezama, John Thistle, and Ufuk Topcu. Furthermore, a special evening session was organized on challenges in supervisory control. Besides the aim of joining and integrating different tracks of research in the area, an important objective of the seminar was to bring together (at least some) members of two large research communities in the area, namely the community of automata, logic, and games in Europe and the U.S. research network EXCAPE (Expeditions in Computer Augmented Program Engineering).

During the seminar, several small circles of participants started or continued joint work. As a general result of the seminar, confirmed by many positive and even enthusiastic comments of participants after the seminar, one may say that a much better understanding and appreciation between the various research branches was established. As one of the participants put it, the seminar was "eye-opening".

As an overview of the areas covered in the talks, we give a short description of the topics studied in the seminar.

Multi-player games. The study of multi-player games is an important extension of the two-player setting. In terms of theoretical study it gives rise to a rich class of questions related to different notions of equilibria, studying computability and complexity results for them, as well as different logics to express them. In terms of practical applicability, various notions of synthesis such as rational synthesis, secure equilibria, assume-guarantee synthesis, assume-admissible synthesis, etc. have been developed to apply the results of multi-player games for synthesis of component-based systems.

Partial-observation games. Partial-observation games extend perfect-information games where players do not have perfect knowledge about the game. This is particularly relevant in control theory, where the controller does not have access to private variables of the plant. The results of partial-observation games have been recently extended to the stochastic setting, as well as for finite-memory strategies, leading to a new framework which can potentially solve interesting applications from the control domain.

Quantitative game models. These are a prominent class of game models for applications to verification and synthesis. In particular, taking real-time constraints into account is especially important for such applications. Timed automata and timed games have already played an important role, as they are a convenient and expressive model enjoying efficient algorithms. Statistical model checking in particular offers a very effective technique for strategy optimization. Robustness analysis of timed models makes the verification process even more faithful.

- Weighted timed games extend timed games with the ability of modeling other quantitative aspects of cyber-physical systems. While the expressive power is greatly improved, the verification and synthesis problems get much more complex than for plain timed automata. Still, algorithms sometimes exist for *approximating* the optimal cost, which in most practical situations is sufficient.
- Timed automata have now reached maturity. Powerful data-structures and efficient symbolic algorithms have been designed to develop efficient algorithms. Statistical model checking is now also used in tools for efficiently optimizing strategies. These tools can now be applied on real-life scenarios, e.g. in home automation and motion planning.
- Probabilistic models form another important class of models of particular interest for representing and reasoning about e.g. systems involving stochastic behaviors. Efficient algorithms have recently been developed for diagnosing probabilistic automata, or for synthesizing strategies that guarantee good performance level with sufficient probability in Markov Decision Processes.

In the quantitative setting, the range of objectives is large; there are mean-payoff objectives, energy objectives, mean-payoff of energy objectives, their Boolean combinations, and combinations of quantitative (e.g., stochastic) semantics with adversarial semantics.

Other domains. Theoretical results developed for games have been generalized to problems in other settings as well. A prime example is that the lower bound for strategy improvement algorithms for parity games was modified to obtain lower bounds for linear-programming solutions, and recent results show that exploiting structures of Markov decision processes it can be established that several classical rules for linear-programming algorithms solve PSPACE-complete problems.

Control engineering. In the field of control engineering, research on supervisory control of discrete event systems and on formal methods in feedback control has recently emphasized distributed and decentralized control architectures that more accurately capture the physical constraints arising in cyber-physical and networked control systems. In these architectures, a set of controllers, possibly with different run-time information about the system, cooperate as a team in order to achieve a specification (either qualitative or quantitative objective) on the entire system behavior, in the presence of a reactive environment. For instance, costly sensors and actuators, as well as costly communication, lead to challenging synthesis problems, both conceptually (e.g., characterization of the information structure) and computationally (e.g., distributed synthesis of the controllers). Another important consideration is to ensure robustness of the synthesized implementation with respect to classes of disturbances on the controlled system.

Researchers are still trying to establish the precise boundary between decidable and undecidable problems in this research domain. It is known that synthesis for both safety and a form of liveness termed non-blockingness, well-understood in a centralized-information setting, becomes undecidable in a decentralized-information setting. But the decidability of special classes of this problem is still an open issue. Establishing concrete bridges between the theory of partial-observation games and such decentralized/distributed control problems for discrete abstractions of cyber-physical systems is an important research issue, both in terms of answering open undecidability questions and in terms of developing efficient synthesis procedures for decidable problems. Similarly, problems of intrusion by malicious agents into control architectures (e.g., taking over actuators or sensors) also lead to new classes of problems where the theory of games with quantitative objectives can be leveraged.

Recently-developed synthesis techniques for "correct-by-construction" controllers in engineering systems have exploited game formulations between the set of controllers on the one hand and the system/environment on the other hand. Related approaches have considered synthesis of the "complete" controller implementation from a "partial" implementation and a sample set of desired behaviors in the reactive environment under consideration.

7.9 Domain-Specific Languages

Organizers: Sebastian Erdweg, Martin Erwig, Richard F. Paige, and Eelco Visser

Seminar No. 15062

Date: February 1–6, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.26

Participants: Lennart Augustsson, Hassan Chafi, William R. Cook, Sebastian Erdweg, Martin Erwig, Matthew Flatt, Andrew Gill, Daco Harkes, Görel Hedin, Steven Kelly, Oleg Kiselyov, Dimitris Kolovos, Shriram Krishnamurthi, Ralf Lämmel, Sonja Maier, Peter D. Mosses, Bruno C. d. S. Oliveira, Klaus Ostermann, Richard F. Paige, Tillmann Rendel, Julia Rubin, Ulrik Pagh Schultz, Yannis Smaragdakis, Friedrich Steimann, Laurence Tratt, Tijs van der Storm, Dániel Varró, Eelco Visser, Markus Völter, Guido Wachsmuth, Eric Walkingshaw

Software systems are the engines of modern information society. Our ability to cope with the increasing complexity of software systems is limited by the programming languages we use to build them. Domain-specific languages (DSLs) successfully address this challenge through linguistic abstraction by providing notation, analysis, verification, optimization, and tooling that are specialized to an application domain. DSLs are already ubiquitous in industrial software development with prominent examples such as HTML, SQL, Make, AppleScript, Matlab, or Simulink.

There is a wide range of methods and techniques for the development of DSLs. Each of these makes different trade-offs that enable different usage scenarios. After the initial design of a DSL, switching to another approach can be very expensive or even impossible. Therefore, the trade-offs and implications of different approaches must be well understood by practitioners from the beginning. However, there is no clear account of what exactly these trade-offs are; neither in industry nor in academia.

The goal of the proposed seminar was to bring together key representatives from the communities that address DSLs from different perspectives: (1) internal DSLs, (2) external DSLs, (3) domain-specific modeling, (4) extensible languages, (5) graph-based languages, and (6) formal semantics. To enable constructive exchange between seminar participants from different communities, the seminar started with one introductory talk per community by a representative. These introductory talks were essential for raising awareness for each other's discipline, the challenges involved, and the problems already solved.

The first day of the seminar was concluded with a poster session. Before the seminar, the organizers invited each participant to prepare and bring a poster that describes their position with respect to the seminar topic. Many participants followed this invitation or used a flip chart for an impromptu presentation. During the poster session, the participants alternated between presenting their own poster and receiving introductions by others. While the seminar did not feature a separate round of introductions at the beginning of the first day, this did not at all hinder discussion and interaction during the talks prior to the poster session. The organizers of this seminar would like to encourage other organizers to consider a poster session as replacement for an introduction round.

After the community and personal introductions on the first day, the second day featured four talks about the "design history" of four existing DSLs. The presenters reported on how the design of their DSLs began, what features turned out to be good, what features turned out to require revision, and how modifications of the design were formed, decided, and implemented. Beyond reporting on their experience, the four talks provided concrete examples of DSLs that could be referred to by all participants during the remainder of the seminar. Subsequently to the design histories, the seminar featured a session on DSL evaluation followed by an industrial panel on industrial DSL requirements.

In the morning of the third day, the participants had the chance to present their latest research results in lightning talks. These were the only talks during the seminar without precise instructions by the seminar organizers. In total, there were eight lightning talks. We observed a high degree of interaction across communities. In the afternoon most participants joined for the excursion: A hike around Schloss Dagstuhl.

Thursday morning was reserved for four talks on DSL type systems. The four talks illustrated different ways of addressing DSL type systems. From a distinguished metalanguage and to automated mechanization to type-system embedding and attribute grammars. The presented work was not mature enough to allow for a meaningful discussion of benefits and disadvantages of the individual approaches. On Thursday afternoon the participants split into two breakout groups on Language Design Patterns and Name Binding. Some participants of the breakout groups decided

to continue exchange and discussion after the seminar. The breakout groups were followed by tool demonstrations, where participants could freely move between demos.

Finally, on Friday morning the seminar ended with a session on establishing a research agenda, that is, relevant research questions that should be addressed by the DSL community. Moreover, the participants found that no new dedicated venue for DSLs needs to be established, because there are sufficiently many venues for DSL research available already.

This report collects the abstracts of the talks, and summarises other activities (including a panel and a discussion on a research agenda). The summaries and abstracts suggest outcomes and potential directions for future scientific research.

7.10 Formal Foundations for Networking

Organizers: Nikolaj Bjørner, Nate Foster, Philip Brighten Godfrey, and Pamela Zave

Seminar No. 15071

Date: Feburary 8–13, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.44

Participants: Aditya Akella, Nikolaj Bjørner, Achim D. Brucker, Marco Canini, Pavol Cerny, Swarat Chaudhuri, Evgeny Chemeritskiy, Shiu-Kai Chin, Bryan Ford, Nate Foster, Vijay Ganesh, Aaron Gember-Jacobson, Philip Brighten Godfrey, Arjun Guha, Arie Gurfinkel, Karthick Jayaraman, Limin Jia, Ethan Katz-Bassett, Shriram Krishnamurthi, Ori Lahav, Nuno Lopes, Ratul Mahajan, Tim Nelson, Aurojit Panda, Panagiotis Papadimitriou, Mark Reitblatt, Jennifer Rexford, Timothy Roscoe, David Rosenblum, Kristin Yvonne Rozier, Andrey Rybalchenko, Mooly Sagiv, Cole Schlesinger, Stefan Schmid, Sharon Shoham Buchbinder, Robert Soulé, David Walker, Alexander L. Wolf, Burkhart Wolff, Yifei Yuan, Vladimir Zakharov, Pamela Zave

The scale and complexity of computer networks has increased dramatically in recent years, driven by the growth of mobile devices, data centers, and cloud computing; increased concerns about security; and generally more widespread and diverse uses of the Internet. Building and operating a network has become a difficult task, even for the most technologically sophisticated organizations.

To address these needs, the research community has started to develop tools for managing this complexity using programming language and formal methods techniques. These tools use domain-specific languages, temporal logics, satisfiability modulo theories solvers, model checkers, proof assistants, software synthesis, etc. to specify and verify network programs.

Yet despite their importance, tools for programming and reasoning about networks are still in a state of infancy. The programming models supported by major hardware vendors require configurations to be encoded in terms of low-level constructs – e.g., hardware forwarding rules and IP address prefixes. To express richer policies, network operators must incorporate "tribal knowledge" capturing requirements that cut across different customers, service-level agreements, and protocols and can easily lead to contradictions. In addition, networks are rarely static, so operators must deal with updates to configurations and the complications that arise during periods of transition or when unexpected failures occur.

The goal of this seminar was to bring together leading practitioners from the areas of formal methods, networking, programming languages, and security, to exchange ideas about problems and solutions, and begin the task of developing formal foundations for networks. The seminar program was grouped into broad categories addressing the following issues:

- Networking Applications (Akella, Gember-Jacobson, Jayaraman, Rexford). What are the key concerns in enterprise, data center, and wide-area networks today? What kinds of modeling, verification, and property-checking tools are operators deploying? What kinds of scalability challenges are they facing?
- Emerging Areas (Papadimitriou, Rozier). What are the key issues in emerging areas such as crowd-sourced networks and aerospace engineering? Can existing tools be easily adapted to these areas? How can new researchers get involved?
- Distributed Systems (Canini, Cerny). What are some techniques for handling the distributed systems issues that arise in modeling and reasoning about networks? How can we exploit these insights to build practical tools for verifying properties in the presence of replicated state, asynchronous communication, and unexpected failures?
- Domain-Specific Tools (Chemeritskiy, Mahajan, Panda, Rybalchenko, Sagiv). What are the best approaches for verifying properties of real-world networks? How can we incorporate features such as dynamic control programs and mutable state? How can we make these tools scale to networks of realistic size?
- General Tools (Brucker, Ganesh, Guha, Jia, Nelson, Rosenblum, Rybalchenko). There is a rich literature on temporal logics, satisfiability modulo theories checkers, model checkers, proof assistants, Datalog, etc. What are the key differences between these tools and how can they be applied to networks?
- Platforms and Models (Guha, Schlesinger, Reitblatt, Walker, Zave). What is the state-of-the-art in network programming? How can we build compilers and hypervisors that correctly translate from high-level models to low-level implementations?
- Program Synthesis (Buchbinder, Chaudhuri, Cerny, Yuan). Synthesis is a promising approach to building correct software, since programs are generated automatically using a verification tool. What are the best current techniques

for using model checkers and satisfiability-modulo theories solvers to generate network configurations, update protocols, and policies?

The seminar comprised four and a half days of presentations, interspersed with discussions, tool demonstrations, and working groups. This report collects the abstracts of the presentations, gives summaries of discussions and working groups, and lists open questions.

7.11 Distributed Cloud Computing

Organizers: Yvonne Coady, James Kempf, Rick McGeer, and Stefan Schmid

Seminar No. 15072

Date: February 8–11, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.64

Participants: Mark Berman, Annette Bieniusa, Justin Cappos, Yvonne Coady, Lars Eggert, Johan Eker, Chip Elliott, Erik Elmroth, Patrick Eugster, Hannes Hartenstein, Oliver Hohlfeld, James Kempf, Igor Konnov, Rick McGeer, Christine Morin, Jörg Ott, Jonathan Pastor, Vivien Quéma, Stefan Schmid, Marc Shapiro, Hagen Wösner, Tim Wood

The Dagstuhl Seminar on Distributed Cloud Computing was held Feb. 8–11, 2015. 22 researchers attended the multidisciplinary seminar from the areas of networking, cloud computing, distributed systems, operations research, security, and system administration. In contrast with the centralized cloud deployment model where applications are restricted to a single mega-data center at some network distance from the customers, in the distributed cloud deployment model, many smaller data centers are deployed closer to customers to supplement or augment the larger mega-data centers, and the smaller data centers are managed as one pooled resource. Two administrative models of a distributed cloud are common today: the integrated model where a single administrative entity controls all the data centers and the federated model where multiple administrative entities control the data centers and users authenticate for resource access using a federated identity management system. Over the course of the 3 day seminar, 15 presentations were given on various aspects of distributed cloud or the disciplinary areas relevant to distributed cloud. The seminar shared two talks with the concurrent seminar on Foundations of Networking and attended one of the Foundations of Networking Talks. Taking the presentations as input, the seminar then broke into three groups to discuss a research agenda for distributed cloud. The groups were requested to come up with 3 questions in their particular area (distributed systems, programming models, and cloud) and two for the other two groups. At the end of the seminar, the group discussed forming a research community around distributed cloud with an annual conference. Currently, a workshop on distributed cloud is held annually, called DCC (for Distributed Cloud Computing). This year's workshop will be held in conjunction with SIGMETRICS in Portland, Oregon in June. Slides, abstracts of the talks and reports from the breakout groups are available in the Dagstuhl content management web site. An extended version of this report appeared in the April 2015 issue of ACM SIGCOMM Computer Communication Review [1].

References

1 Yvonne Coady, Oliver Hohlfeld, James Kempf, Rick McGeer, and Stefan Schmid. Distributed cloud computing: Applications, status quo, and challenges. *SIGCOMM Comput. Commun. Rev.*, 45(2):38–43, April 2015. DOI: 10.1145/2766330.2766337

Fig. 7.4
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.12 Holistic Scene Understanding

Organizers: Jiří Matas, Vittorio Murino, and Bodo Rosenhahn

Seminar No. 15081

Date: February 15–20, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.80

Participants: Gabriel Brostow, Marco Cristani, Alessio Del Bue, Markus Enzweiler, Michele Fenzi, Sanja Fidler, Bob Fisher, Jan-Michael Frahm, Jürgen Gall, Shaogang Gong, Abhinav Gupta, Michal Havlena, Roberto Henschel, Esther Horbert, Jörn Jachalsky, Ron Kimmel, Reinhard Klette, Laura Leal-Taixé, Oisin Mac Aodha, Jiri Matas, Greg Mori, Vittorio Murino, Caroline Pantofaru, Matthias Reso, Anna Rohrbach, Bodo Rosenhahn, Bernt Schiele, Konrad Schindler, Min Sun, Raquel Urtasun, Sebastiano Vascon, Stefan Walk, Jan Dirk Wegner, Michael Yang, Angela Yao

Motivations

To *understand* a scene in a given image or video is much more than to *simply* record and store it, extract some features and eventually recognize an object. The overall goal is to find a mapping to derive semantic information from sensor data. Purposive Scene understanding may require a different representation for different specific tasks. The task itself can be used as prior but we still require an in-depth understanding and balancing between local, global and dynamic aspects which can occur within a scene. For example, an observer might be interested to understand from an image if there is a person present or not, and beyond that, if it is possible to look for more information, e.g. if the person is sitting, walking or raising a hand, etc.

When people move in a scene, the specific time (e.g. 7:30 in the morning, workdays, weekend), the weather (e.g. rain), objects (cars, a bus approaching a bus stop, crossing bikes, etc.) or surrounding people (crowded, fast moving people) yield to a mixture of low-level and high-level, as well as abstract cues, which need to be jointly analyzed to get an in-depth understanding of a scene. In other words, generally speaking, the so-called *context* is to be considered for a comprehensive scene understanding, but this information, while it is easily captured by human beings, is still difficult to obtain from a machine.

Holistic scene interpretation is crucial to design the next generation of recognition systems, which are important for several applications, e.g. driver assistance, city modeling and reconstruction, outdoor motion capture and surveillance.

With such topics in mind, the aim of this workshop was to discuss which are the sufficient and necessary elements for a complete scene understanding, i.e. what it really means to *understand* a scene. Specifically, in this workshop, we wanted to explore methods that are capable of representing a scene at different level of semantic granularity and modeling various degrees of interactions between objects, humans and 3D space. For instance, a scene-object interaction describes the way a scene type (e.g. a dining room or a bedroom) influences objects' presence, and vice versa. An object-3D-layout or human-3D-layout interaction describes the way the 3D layout (e.g. the 3D configuration of walls, floor and observer's pose) biases the placement of objects or humans in the image, and vice versa. An object-object or object-human interaction describes the way objects, humans and their pose affect each other (e.g. a dining table suggests that a set of chairs are to be found around it). In other words, the 3D configuration of the environment and the relative placements and poses of the objects and humans therein, the associated dynamics (relative distance, human body posture and gesture, gazing, etc.), as well as other contextual information (e.g., weather, temperature, etc.) support the holistic understanding of the observed scene.

As part of a larger system, understanding a scene semantically and functionally allows to make predictions about the presence and locations of unseen objects within the space, and thus predict behaviors and activities that are yet to be observed. Combining predictions at multiple levels into a global estimate can improve each individual prediction.

Since most scenes involve humans, we were also interested in discussing novel methods for analyzing group activities and human interactions at different levels of spatial and semantic resolution. As advocated in recent literature, it is beneficial to solve the problem of tracking individuals and understand their activities in a joint fashion by combining bottom-up evidence with top-down reasoning as opposed to attack these two problems in isolation.

Top-down constraints can provide critical contextual information for establishing accurate associations between detections across frames and, thus, for obtaining more robust tracking results. Bottom-up evidence can percolate upwards so as to automatically infer action labels for determining activities of

individual actors, interactions among individuals and complex group activities. But of course there is more than this, it is indeed the cooperation of both data flows that makes the inference more manageable and reliable in order to improve the comprehension of a scene.

We gathered researchers which are not only well-known in Computer Vision areas such as object detection, classification, motion segmentation, crowd and group behavior analysis or 3D scene reconstruction, but also Computer Vision affiliated people from other communities in order to share each others point of view on the common topic of scene understanding.

Goals

Our main goals of the seminar can be summarized as follows:

- Address holistic scene understanding, a topic that has not been discussed before in detail at previous seminars, with special focus on a multidisciplinary perspective for sharing or competing the different views.
- Gather well-known researchers from the Computer Vision, Machine Learning, Social Sciences (e.g. Cognitive Psychology), Neuroscience, Robotics and Computer Graphics communities to compare approaches to representing scene geometry, dynamics, constraints as well as problems and task formulations adopted in these fields. The interdisciplinary scientific exchange is likely to enrich the communities involved.
- Create a platform for discussing and bridging topics like perception, detection, tracking, activity recognition, multi--people multi-object interaction and human motion analysis, which are surprisingly treated independently in the communities.
- Publication of an LNCS post-proceedings as previously done for the 2006, 2008 and 2010 seminars. These will include the scientific contributions of participants of the Seminar, focusing specially on the discussed topics presented at the Seminar.

Organization of the seminar

During the workshop we discussed different modeling techniques and experiences researchers have collected. We discussed sensitivity, time performance and e.g. numbers of parameters required for special algorithms and the possibilities for context-aware adaptive and interacting algorithms. Furthermore, we had extensive discussions on open questions in these fields.

On the first day, the organizers provided general information about Dagstuhl seminars, the philosophy behind Dagstuhl and the expectations to the participants. We also clarified the kitchen-rules and organized a running-group for the early mornings (5 people participated frequently!).

Social event. On Wednesday afternoon we organized two afternoon event: One group made a trip to Trier, and another group went on a 3h hike in the environment.

Working Groups. To strongly encourage discussions during the seminar, we organized a set of working groups on the first day (with size between 8–12 people). As topics we selected

- What does “Scene Understanding” mean?
- Dynamic Scene: Humans.
- Recognition in static scenes (in 3D).

There were two afternoon slots reserved for these working groups and the outcome of the working groups has been presented in the Friday morning session.

LNCS Post-Proceedings. We will edit a Post-Proceeding and invite participants to submit articles. In contrast to standard conference articles, we allow for more space (typically 25 single-column pages) and allow to integrate open questions or preliminary results, ideas, etc. from the seminar into the proceedings. Additionally, we will enforce joint publications of participants who started to collaborate after the seminar. All articles will be reviewed by at least two reviewers and based on the evaluation, accepted papers will be published. We will publish the proceeding at the Lecture Notes in Computer Science (LNCS-Series) by Springer. The papers will be collected during the summer months.

Overall, it was a great seminar and we received very positive feedback from the participants. We would like to thank castle Dagstuhl for hosting the event and are looking forward to revisit Dagstuhl whenever possible.

7.13 Limitations of Convex Programming: Lower Bounds on Extended Formulations and Factorization Ranks

Organizers: Hartmut Klauck, Troy Lee, Dirk Oliver Theis, and Rekha R. Thomas

Seminar No. 15082

Date: February 15–20, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.109

Participants: Alexander Barvinok, LeRoy B. Beasley, Greg Blekherman, Daniel Dadush, Ronald de Wolf, Samuel Fiorini, Nicolas Gillis, Francois Glineur, João Gouveia, Sebastian Gruler, Alexander Guterman, Volker Kaibel, Hartmut Klauck, Kaie Kubjas, James R. Lee, Troy Lee, Arnau Padrol Sureda, Kanstantsin Pashkovich, Teresa Piovesan, Sebastian Pokutta, Raman Sanyal, Markus Schweighofer, Yaroslav Shitov, David Steurer, Jonathan Swenson, Dirk Oliver Theis, Rekha R. Thomas, Hans Raj Tiwary, Antonios Varvitsiotis, Stefan Weltge

The topic of this seminar was the rapidly developing notion of *cone rank* of a matrix/polytope that is an important invariant controlling several properties of the matrix/polytope with connections to optimization, communication complexity and theoretical computer science. This meeting was a follow-up to the 2013 Dagstuhl seminar *13082: "Communication Complexity, Linear Optimization, and lower bounds for the nonnegative rank of matrices"* organized by Leroy Beasley, Hartmut Klauck, Troy Lee and Dirk Oliver Theis.

The cone rank of a nonnegative matrix is an ordered notion of matrix rank with emerging applications in several fields. A well-known example is the nonnegative rank of a nonnegative matrix that appears in areas ranging from communication complexity, to statistics, to combinatorial optimization, and algebraic complexity theory. A related notion arising as an invariant in representations of convex sets as projections of affine slices of the positive semidefinite (psd) cone is the positive semidefinite (psd) rank. The psd rank is a very new quantity and is still relatively poorly understood.

The purpose of this seminar was to bring together researchers from optimization, computer science, real/convex/tropical algebraic geometry, and matrix theory, to discuss relevant techniques from each area that can contribute to the development of both the theory and computation of cone ranks and cone factorizations of nonnegative matrices, as well as their many emerging applications.

In optimization and computer science, a common approach to finding approximate solutions to NP-hard problems is to look at tractable convex relaxations of the problem as either a linear or semidefinite program. An optimal solution to such a relaxation gives a bound on the objective value of the original problem. While much previous work has focused on specific relaxations of a problem, or a family of relaxations coming from a hierarchy, cone ranks allow the study of the best possible linear, semidefinite, or other convex formulations of a NP-hard problem independent of specific construction methods. These formulations all write the underlying feasible set as the projection of an affine slice of a closed convex cone and is commonly refereed to as an extended formulation of the underlying feasible region. The nonnegative rank of a polytope is the smallest size of a linear extended formulation of the polytope while psd rank of the polytope is the size of the smallest possible semidefinite extended formulation of the polytope. Linear extended formulations are the best understood so far and an exciting development in this area is the recent breakthrough by Rothvoß showing that the matching polytope does not admit a polynomial sized linear extended formulation, settling a notorious open problem in combinatorial optimization. Very recently, there has also been exciting new developments in the area of psd rank such as the result of Lee, Raghavendra, and Steurer that shows that the psd rank of certain polytopes such as the traveling salesman polytope of a graph with n vertices must be exponential in n.

In the field of communication complexity, nonnegative and psd ranks are exactly characterized by a model of randomized and quantum communication complexity, respectively. This connection has allowed tools from communication and information theory to help create lower bounds for these ranks.

A central question in the field of real algebraic geometry is the semidefinite representability of convex sets. While polytopes only project to polytopes, affine slices of psd cones have much greater expressive power as their projections are convex sets, which allows the definition of psd rank for semi algebraic convex sets. Psd rank has inherent semi algebraic structure and its study crosses over into real and convex algebraic geometry, algebraic complexity, and semidefinite programming.

Fig. 7.5
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of (left) Dr. Christian Lindig and (right) TreeState Productions GmbH, www.treestate.de.

7.14 Smart Buildings and Smart Grids

Organizers: Hans-Arno Jacobsen, Randy H. Katz, Hartmut Schmeck, and Christoph Goebel

Seminar No. 15091

Date: February 22–27, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.2.128

Participants: Florian Allerding, Birger Becker, Bert Claessens, Hermann de Meer, Victor del Razo, Christoph Doblander, Frank Eliassen, Nicolas Gast, Christoph Goebel, Manuel Görtz, Kai Heussen, Hanno Hildmann, Longbo Huang, Hans-Arno Jacobsen, Randy H. Katz, Srinivasan Keshav, Jean-Yves Le Boudec, Florian Michahelles, Milan Milenkovic, Johanna Myrzik, Peter Noglik, Mario Paolone, Anthony Papavasiliou, Yvonne-Anne Pignolet, Jose Rivera, Jiri Rojicek, Hartmut Schmeck, Mischa Schmidt, Hans-Peter Schwefel, Joachim Sokol, Thorsten Staake, Kai Strunz, Verena Tiefenbeck, Anwar Ul Haq, Andreas Veit, Holger Ziekow

Motivation

Motivated by the increasing importance of producing and consuming energy more sustainably, a new and highly dynamic research community within computer science has evolved: Energy Informatics (EI). Researchers active in the EI field investigate information age solutions for monitoring and controlling large cyber-physical infrastructures with a focus on the following goals: (i) an overall reduction of the energy consumption of these infrastructures, and (ii) the integration of distributed renewable energy sources into these infrastructures. This seminar focused on two use cases of existing cyber-physical systems, buildings and power grids. These use cases were chosen due to their relevance in terms of energy footprint. The seminar has three major goals: (i) to provide a forum for leading EI researchers to discuss their recent research on Smart Buildings and Smart Grids, (ii) to further elaborate EI research agenda and methods, and (iii) to kick-start new research projects with industry.

Smart Buildings. Modern buildings already incorporate increasingly sophisticated Building Management Systems (BMS) that integrate building control with improved sensors and better data collection and presentation capabilities. However, these systems currently only enable simple, decoupled control of building services like lighting, ventilation, heating and cooling. Their architectures and Application Programming Interfaces (APIs) are not standardized, and often proprietary: only the BMS vendor can add functionality. This slows the pace of innovation, buildings remain rigid in the functions and services they provide, and their quality and effectiveness remain difficult to quantify. Contemporary BMS attempt to achieve global service levels based on local control instead of meeting individual occupant requirements based on global control. Standardized building management APIs and scalable middleware solutions that enable reliable communication between building sensors, users, control systems, and machinery could accelerate energy innovation in the building sector.

Smart Grids. Contemporary electricity grids and markets were designed for a scenario in which large and mostly fossil-fueled power plants are dispatched to meet an almost inflexible demand. Achieving sustainable energy supply, however, requires moving towards a scenario where the variable power supplied by distributed renewable resources like wind and solar has to be absorbed by supply-following loads and energy storage whenever it is available. Thus, instead of dispatching a relatively small number of large generators, the large-scale integration of new types of generators and loads into electric grids requires new types of information systems for monitoring and controlling them, while making efficient use of existing assets. The task of controlling large numbers of flexible loads, e.g., air conditioning systems in buildings, electric vehicles, and small-scale energy storage systems, while guaranteeing overall system stability, is highly demanding in terms of computational complexity, required data communication and data storage. In the Smart Grid space, the challenge faced by EI researchers is to develop and carefully evaluate new ideas and actual system components enabling Smart Grid systems that are scalable, efficient, reliable, and secure.

Organization of the Seminar

The week-long workshop plan was as follows. Day 1 introduced the attendees to each other, and set the stage through invited tutorial presentations and brainstorming sessions. Day 2 was spent in breakouts focused on identifying the research challenges and opportunities, organized by application area such as Smart Buildings or Energy Grids, based on attendee interest and expertise. On Day 2, we also held the first out of two presentation session, where participants could give a short overview about their current research. Day 3 was used to assess the

workshop at mid-stream, conduct group discussions, and make necessary corrections. Initial writing assignments, to document the discussions of the breakout sessions, were made on this day, as well. Day 4 consisted of a second round of breakouts focusing on enablers and crosscutting issues (e.g., data management, system design patterns, and human machine interaction) and the second participants' presentation session. Work on completing the report draft continued on that day. The last day consisted of the reviewing of the report draft, and through group discussion, identify the summary findings and recommendations.

7.15 Bridging Information Visualization with Machine Learning

Organizers: Daniel A. Keim, Tamara Munzner, Fabrice Rossi, and Michel Verleysen

Seminar No. 15101

Date: March 1–6, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.1

Participants: Daniel Archambault, Francois Blayo, Kerstin Bunte, Miguel Á. Carreira-Perpiñán, Ignacio Díaz Blanco, David S. Ebert, Alex Endert, Thomas Ertl, Barbara Hammer, Helwig Hauser, Stephen Ingram, Samuel Kaski, Daniel A. Keim, Bongshin Lee, John A. Lee, Torsten Möller, Bassam Mokbel, Tamara Munzner, Ian Nabney, Stephen North, Eli Parviainen, Fernando Paulovich, Jaakko Peltonen, William Ribarsky, Fabrice Rossi, Frank-Michael Schleif, Michael Sedlmair, Cagatay Turkay, Jarke J. van Wijk, Michel Verleysen, Thomas Villmann, Daniel Weiskopf, William Wong, Jing Yang, Leishi Zhang, Blaz Zupan

Motivations and context of the seminar

Following the success of Dagstuhl seminar 12081 "Information Visualization, Visual Data Mining and Machine Learning" [1, 2], which provided to the participants from the IV and ML communities the ground for understanding each other, this Dagstuhl seminar aimed at bringing once again the visualization and machine learning communities together.

Information visualization and visual data mining leverage the human visual system to provide insight and understanding of unorganized data. Visualizing data in a way that is appropriate for the user's needs proves essential in a number of situations: getting insights about data before a further more quantitative analysis (e.g., for expert selection of a number of clusters in a data set), presenting data to a user through well-chosen table, graph or other structured representations, relying on the cognitive skills of humans to show them extended information in a compact way, etc.

The scalability of visualization methods is an issue: human vision is intrinsically limited to between two and three dimensions, and the human preattentive system cannot handle more than a few combined features. In addition the computational burden of many visualization methods is too large for real time interactive use with large datasets. In order to address these scalability issues and to enable visual data mining of massive sets of high dimensional data (or so-called "big data"), simplification methods are needed, so as to select and/or summarize important dimensions and/or objects.

Traditionally, two scientific communities developed tools to address these problems: the machine learning (ML) and information visualization (IV) communities. On the one hand, ML provides a collection of automated data summarizing/compression solutions. Clustering algorithms summarize a set of objects with a smaller set of prototypes, while projection algorithms reduce the dimensionality of objects described by high-dimensional vectors. On the other hand, the IV community has developed user-centric and interactive methods to handle the human vision scalability issue.

Building upon seminar 12081, the present seminar aimed at understanding key challenges such as interactivity, quality assessment, platforms and software, and others.

Organization

The seminar was organized in order to maximize discussion time and in a way that avoided a conference like program with classical scheduled talks. After some lightning introduction by each participant, the seminar began with two tutorial talks one about machine learning (focused on visualization related topics) followed by another one about information visualization. Indeed, while some attendants of the present seminar participated to seminar 12081, most of the participants did not. The tutorials helped establishing some common vocabulary and giving an idea of ongoing research in ML and IV.

After those talks, the seminar was organized in parallel working groups with periodic plenary meeting and discussions, as described below.

Topics and groups

After the two tutorials, the participants spend some time identifying topics they would like to discuss during the seminar. Twenty one emerged:

1. Definition and analysis of quantitative evaluation measures for dimensionality reduction (DR) methods (and for other methods);
2. In the context of dimensionality reduction: visualization of quality measures and of the sensitivity of some results to user inputs;

3. What IV tasks (in addition to DR related tasks) could benefit from ML? What ML tasks could benefit from IV?
4. Reproducible/stable methods and the link of those aspects to sensitivity and consensus results;
5. Understanding the role of the user in mixed systems (which include both a ML and an IV component);
6. Interactive steerable ML methods (relation to intermediate results);
7. Methods from both fields for dynamic multivariate networks;
8. ML methods that can scale up to IV demands (especially in terms of interactivity);
9. Interpretable/transparent decisions;
10. Uncertainty;
11. Matching vocabularies/taxonomies between ML and IV;
12. Limits to ML;
13. Causality;
14. User guidance: precalculating results, understanding user intentions;
15. Mixing user and data driven evaluation (leveraging a ROC curve, for instance);
16. Privacy;
17. Applications and use cases;
18. Prior knowledge integration;
19. Formalizing task definition;
20. Usability;
21. Larger scope ML.

After some clustering and voting those topics were merged into six popular broader subjects which were discussed in working groups through the rest of the week:

1. Dynamic networks
2. Quality
3. Emerging tasks
4. Role of the user
5. Reproducibility and interpretability
6. New techniques for Big Data

The rest of the seminar was organized as a series of meeting in working groups interleaved with plenary meetings which allowed working groups to report on their joint work, to steer the global process, etc.

Conclusion

As reported in the rest of this document, the working groups were very productive as was the whole week. In particular, the participants have identified a number of issues that mostly revolve around complex systems that are being built for visual analytics. Those systems need to be scalable, they need to support rich interaction, steering, objective evaluation, etc. The results must be stable and interpretable, but the system must also be able to include uncertainty into the process (in addition to prior knowledge). Position papers and roadmaps have been written as a concrete output of the discussions on those complex visual analytics systems.

The productivity of the week has confirmed that researchers from information visualization and from machine learning share some common medium to long term research goals. It appeared also clearly that there is still a strong need for a better understanding between the two communities. As such, it was decided to work on joint tutorial proposals for upcoming IV and ML conferences. In order to facilitate the exchange between the communities outside of the perfect conditions provided by Dagstuhl, the blog "Visualization meets Machine Learning[39]" was initiated.

It should be noted finally that the seminar was very appreciated by the participants as reported by the survey. Because of the practical organization of the seminar, participants did not know each other fields very well and it might have been better to allows slightly more time for personal introduction. Some open research questions from each field that seems interesting to the other fields could also have been presented. But the positive consequences of avoiding a conference like schedule was very appreciated. The participants were pleased by the ample time for discussions, the balance between the two communities and the quality of the discussions. Those aspects are quite unique to Dagstuhl.

References

1 Daniel A. Keim, Fabrice Rossi, Thomas Seidl, Michel Verleysen, and Stefan Wrobel. Dagstuhl Manifesto: Information Visualization, Visual Data Mining and Machine Learning (Dagstuhl Seminar 12081). *Informatik-Spektrum*, 35:58–83, 8 2012.

2 Daniel A. Keim, Fabrice Rossi, Thomas Seidl, Michel Verleysen, and Stefan Wrobel, (editors). *Information Visualization, Visual Data Mining and Machine Learning (Dagstuhl Seminar 12081)*, Dagstuhl Reports, 2(2):58–83, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2012. http://dx.doi.org/10.4230/DagRep.2.2.58

[39] http://vismeetsml.b.uib.no/

7.16 Secure Routing for Future Communication Networks

Organizers: Amir Herzberg, Matthias Hollick, and Adrian Perrig

Seminar No. 15102

Date: March 1–4, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.28

Participants: Steven Bellovin, Saleem Bhatti, Randy Bush, Joel M. Halpern, Amir Herzberg, Matthias Hollick, Ivan Martinovic, Rossella Mattioli, Cristina Nita-Rotaru, Michael Noisternig, Panagiotis Papadimitratos, Adrian Perrig, Raphael Reischuk, Alvaro Retana, Michael Schapira, Thomas C. Schmidt, Jean-Pierre Seifert, Haya Shulman, Mahesh Tripunitara, Gene Tsudik, Laurent Vanbever, Matthias Wählisch

Routing is a fundamental mechanism in communication networks, and its security is critical to ensure availability and to prevent attacks; however, developing and deploying secure routing mechanisms is still a challenge. Routing is the process by which information is passed via the communication network, from source to destination, via a series of intermediary nodes/routers. Routing attacks include route-hijacking, i.e., diverting traffic to an adversary-controlled router, and denial-of-service attacks exploiting the routing mechanism, i.e., preventing communication (in parts or the entire network), e.g., by malicious dropping of packets by a router.

Routing, and even more secure routing, are complex problems with many variants. In particular, the Internet is a federation of many domains (usually referred to as autonomous systems (ASes)), each managed by a separate organization; there are separate standard protocols for routing inside an AS (intra-domain routing) and for routing from a source in one AS to a destination in a different AS (inter-domain routing). Significant efforts are dedicated to securing intra-domain routing protocols and inter-domain routing protocols; in addition, significant efforts are also dedicated to the design of completely new Internet architectures that include secure routing mechanisms.

Another categorization of routing mechanisms and challenges involves mobility. Many routing protocols, including standard Internet routing, are designed for a mostly static topology, where connections between routers are relatively stable. However, communication is increasingly performed among mobile devices. There are many efforts and challenges in the design of (secure) routing mechanisms for highly mobile networks, e.g., between tiny wireless sensors, swarms of tiny robots, or simply mobile users (e.g., upon catastrophic failure to regular infrastructure).

There is also a need to re-evaluate and possibly re-design routing mechanisms and security measures, to address changes in the way the Internet is used, and in the presence of new security challenges. In particular, is there a need to adapt routing to facilitate, and/or take advantage of, cloud services, and to support security for them? Is there a need to adapt routing to the increased threat of Denial-of-Service attacks, or to facilitate widespread provision of Quality-of-Service? Should routing be modified to take into account energy considerations, or to take advantage of and facilitate Software Defined Networking (SDN)? If modifications are made for these goals, how does this affect routing systems' attack surface? Finally, is there a need to modify routing and its security mechanisms, as a result of the recent revelations regarding the scope of abuse of routing by powerful nation-state adversaries?

In summary, to advance routing security in the aforementioned topic areas, a number of significant research problems need to be addressed, and identifying these problems was the goal of this seminar. The first objective was to facilitate brainstorming and exchange of ideas among experts working in different areas and types of secure networking, leading to an improved understanding of the different aspects of secure routing. The second objective was to identify the most important research challenges and to devise a roadmap towards addressing urgent issues. Through the seminar, we aimed at opening up new avenues of research in the area of routing security. For the given focus areas of the seminar, we contributed to the following key research challenges:

- Routing Security by Design for a Future Internet: the challenge was to overcome the limitations and confined models imposed by today's Internet. Both clean slate as well as evolutionary approaches towards a secure-by-design future Internet were discussed.
- Inter-domain Routing Security and Intra-domain Routing Security: challenges addressed in inter-domain routing were the reconciliation of potentially conflicting security interests across multiple domains and resilience against recently pub-

lished attacks. Intra-domain routing is underrepresented in research; here, the seminar aimed at identifying the key research challenges towards a research roadmap.
- Routing Security in Mobile/Wireless Networks, and in Delay- and Disruption-tolerant Networks: the main goal within the seminar was to identify possible ways to provide routing security in light of the severely limited resources and special characteristics of mobile and wireless systems.
- Quality of Service (QoS) and Denial of Service (DoS) aspects of Routing Security: the challenge was to jointly consider security considerations and QoS aspects, both in theory and practice.

To address these challenges, the seminar was organized in six working groups which are described in the full report.

7.17 Computational Geometry

Organizers: Otfried Cheong, Jeff Erickson, and Monique Teillaud

Seminar No. 15111

Date: March 8–13, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.41

Participants: Peyman Afshani, Annamaria Amenta, Franz Aurenhammer, Maike Buchin, Sergio Cabello, Siu-Wing Cheng, Otfried Cheong, Jinhee Chun, Mark de Berg, Vin de Silva, Olivier Devillers, Tamal K. Dey, Michael Gene Dobbins, Anne Driemel, Ioannis Z. Emiris, Jeff Erickson, Jie Gao, Marc Glisse, Leonidas J. Guibas, Michael Kerber, David G. Kirkpatrick, Rolf Klein, Joseph S. B. Mitchell, Wolfgang Mulzer, Elizabeth Munch, Yota Otachi, Jeff M. Phillips, Natan Rubin, Vera Sacristan, Maria Saumell, Ludmila Scharf, André Schulz, Raimund Seidel, Donald Sheehy, Jonathan Shewchuk, Anastasios Sidiropoulos, Fabian Stehn, Monique Teillaud, Csaba Toth, Carola Wenk, Nicola Wolpert

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, robotics, geographic information systems, 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, significant real-world impact, and strong intellectual connections with other computing and mathematics disciplines.

Seminar Topics

The emphasis of the seminar was on presenting recent developments in computational geometry, as well as identifying new challenges, opportunities, and connections to other fields of computing. In addition to the usual broad coverage of emerging results in the field, the seminar included invited survey talks on two broad and overlapping focus areas that cover a wide range of both theoretical and practical issues in geometric computing. Both focus areas have seen exciting recent progress and offer numerous opportunities for further cross-disciplinary impact.

Computational topology and topological data analysis. Over the last decade, computational topology has grown from an important subfield of computational geometry into a mature research area in its own right. Results in this field combine classical mathematical techniques from combinatorial, geometric, and algebraic topology with algorithmic tools from computational geometry and optimization. Key developments in this area include algorithms for modeling and reconstructing surfaces from point-cloud data, algorithms for shape matching and classification, topological graph algorithms, new generalizations of persistent homology, practical techniques for experimental low-dimensional topology, and new fundamental results on the computability and complexity of embedding problems. These results have found a wide range of practical applications in computer graphics, computer vision, robotics, sensor networks, molecular biology, data analysis, and experimental mathematics.

Geometric data analysis. Geometric data sets are being generated at an unprecedented scale from many different sources, including digital video cameras, satellites, sensor networks, and physical simulations. The need to manage, analyze, and visualize dynamic, large-scale, high-dimensional, noisy data has raised significant theoretical and practical challenges not addressed by classical geometric algorithms. Key developments in this area include new computational models for massive, dynamic, and distributed geometric data; new techniques for effective dimensionality reduction; approximation algorithms based on coresets and other sampling techniques; algorithms for noisy and uncertain geometric data; and geometric algorithms for information spaces. Results in this area draw on mathematical tools from statistics, linear algebra, functional analysis, metric geometry, geometric and differential topology, and optimization, and they have found practical applications in spatial databases, clustering, shape matching and analysis, machine learning, computer vision, and scientific visualization.

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

talent to the seniors of the field. They have arguably been the most influential meetings in the field of computational geometry.

The organizers held a *lottery* for the second 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. Researchers on the initial list who were not selected by the lottery were notified by us separately per email, so that they knew that they were not forgotten, and to reassure them that – with better luck – they will have another chance in future seminars. The seminar has now a more balanced attendance in terms of seniority and gender than in the past.

This year, 41 researchers from various countries and continents attended the seminar, showing the strong interest of the community for this event. The feedback from participants was very positive.

No other meeting in our field allows young researchers to meet with, get to know, and work with well-known and senior scholars to the extent possible at the Dagstuhl Seminar.

We warmly thank the scientific, administrative and technical staff at Schloss Dagstuhl! Dagstuhl allows people to really meet and socialize, providing them with a wonderful atmosphere of a unique closed and pleasant environment, which is highly beneficial to interactions. Therefore, Schloss Dagstuhl itself is a great strength of the seminar.

7.18 Network Calculus

Organizers: Florin Ciucu, Markus Fidler, Jörg Liebeherr, and Jens Schmitt

Seminar No. 15112

Date: March 8–11, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.63

Participants: Sami Akin, Hussein Al-Zubaidy, Michael Beck, Nico Becker, Daniel Berger, Steffen Bondorf, Anne Bouillard, Marc Boyer, Peter Buchholz, Almut Burchard, Florin Ciucu, Markus Fidler, Reinhard German, Fabien Geyer, Yashar Ghiassi-Farrokhfal, James Gross, Kai-Steffen Jens Hielscher, Yuming Jiang, Eduard Jorswieck, George Kesidis, Kai Lampka, Jörg Liebeherr, Krishna S. Pandit, Felix Poloczek, Amr Rizk, Adi Rosén, Gabriel Scalosub, Jens Schmitt, Giovanni Stea, Hao Wang, Nataly Youssef

The network calculus has established as a versatile methodology for the queueing analysis of resource sharing based systems. Its prospect is that it can deal with problems that are fundamentally hard for alternative methodologies, based on the fact that it works with bounds rather than striving for exact solutions. The high modelling power of the network calculus has been transposed into several important applications for network engineering problems, traditionally in the Internet's Quality of Service proposals IntServ and DiffServ, and more recently in diverse environments such as wireless networks, sensor networks, switched Ethernets, Systems-on-Chip, as well as smart grids.

The goal of this Dagstuhl seminar was to gather the deterministic and stochastic network calculus community, to discuss recent research activities, to identify future research questions, and to strengthen cooperation. Topics of this Dagstuhl seminar were:

Wireless systems: for the analysis of wireless networks, a question of interest is how the stochastic properties of wireless channels impact delay and backlog performance. The usual statistical models for radio signals in a propagation environment do not lend themselves easily to a queueing model. Promising methods that were elaborated in the seminar are effective capacities and a recent network calculus of fading channels.

Lower bounds and tightness of bounds: based on the ability to solve some fundamentally hard queueing problems, the stochastic network calculus is regarded as a valuable alternative to the classical queueing theory. The derivation of performance bounds in the stochastic network calculus, e.g., for backlog, and delay, frequently exploits well known tail estimates, such as Chernoff bound and others. The tightness of these bounds and alternative more accurate models and techniques, such as Martingale bounds, were a topic of the seminar.

Network topology: a remarkable quality of the network calculus is that it includes a variety of systems that can be composed to arbitrary network topologies. Various analytical as well as numerical approaches have been explored to analyze different types of topologies, such as line topologies or feed-forward networks. The goal of this seminar was to identify relevant classes of topologies, their defining properties, and corresponding methods.

Parallel systems: the area of performance evaluation of parallel systems has recently become increasingly important due to the prevalence of modern parallel computational models. It is thus a great opportunity for the network calculus community to develop new models and methods which can enable a fundamental and broad understanding of the performance of parallel systems. At the seminar, recent approaches to parallel systems have been discussed.

Related methods: the network calculus has a number of rather unexplored and unexploited connections to related methods in the areas of competitive analysis, adversarial queueing theory, and robust queueing theory that may offer a significant potential for future research. At the seminar, researchers from related fields provided valuable new input to the network calculus community.

During the seminar, we discussed and (partly) answered the following questions:

What are the requirements on a wireless network calculus? Given the increasing importance of wireless communications, the seminar featured two sessions comprising seven presentations on wireless systems, where different approaches and their applications were discussed. Subsequently, a wireless roadmap discussion was centered around the following questions:

- How to model wireless channels and systems?
- What are the most relevant future systems and technologies?

- Which assumptions are needed, which can be safely made?
- What kind of results are needed, which theories can provide these?

With regard to the questions above, we highlight some of the main aspects that were elaborated on during the seminar. The methods that were presented include

- effective capacities,
- impairment models (duality with left-over service curves of scheduling),
- (min, x)-calculus for fading channels,
- capacity-delay-error boundaries,
- central limit theorem,
- Martingale bounds.

Providing different pros, a common basis of many of these methods was found to be due to the prevailing use of moment generating functions (Laplace transforms or Mellin transforms). Relevant systems that were discussed are cognitive radio, 3GPP, MIMO, spatial multiplexing, automatic repeat request, and medium access control. Some fundamental aspects of modelling wireless systems are the assumptions that are required today. Typical choices include

- service increments:
  - independent,
  - Markovian, Gilbert-Elliott channel,
- in-order delivery,
- error-free, instantaneous feedback channel,
- instantaneous retransmission of erroneous data,
- channel state information.

During the discussion, the need for transfer domains beyond Gilbert-Elliott models was raised. Also, the introduction of a notion of time into information-theoretic concepts, such as channel capacity, was discussed and finite-block length capacity results were brought up. Topics of further interest included spatial aspects of wireless networks, interference, and multi-hop networks in general. Regarding the solutions that can be obtained, a tradeoff between exactness and analytical closed forms became apparent. In particular, in system optimization analytical solutions were mentioned to be most useful to obtain derivatives of relevant performance measures. The discussion also touched upon some more general aspects such as qualitative vs. quantitative results, where many practical applications may not require exact results but can benefit from measurable rules of thumb.

What are most promising future research topics in the network calculus? This question was elaborated on in group work sessions, where the task was to identify an upcoming, relevant research topic where performance evaluation can be expected to make a key contribution. The discussion was guided by the following questions:

- What are the requirements for theory, which assumptions can be made?
- Which results would be needed from theory?
- How would a model/approach look like?
- What would be the best case outcome?
- Which body of theory could provide such results?
- What would be a good topic/method/approach for a PhD dissertation in this area?

Relevant topics in the network calculus were found to include cross-layer design, industrial communication, systems on chip, networks on chip, data center communication, and big data. A strategic orientation may also focus on new and unorthodox problems such as

- just-in-time manufacturing,
- renewable energy, smart grid,
- caching,
- financial engineering,
- road traffic,

where the intuitive concept of envelopes as used by the network calculus may be beneficial for many applications in industry. Methodological aspects that may pose relevant and interesting challenges were discussed in the areas of:

- re-entrant lines, particularly stability of such systems,
- max-min problems,
- derivative constraints, e.g., in modelling of batteries,
- network topologies, particularly non-feed forward networks.

Making network calculus happen: computational aspects, application modelling, tool support. Clearly, for network calculus to become a standard technique in performance modelling and analysis of networked and distributed systems it is crucial to arrive at computable solutions, demonstrate its strengths in diverse applications and provide software tools to support performance engineers in their daily tasks. As these different issues are interrelated on many levels two sessions with nine presentations were devoted to them. Among the different issues raised during these presentations and the corresponding discussions were the following:

- What are suitable novel application domains for network calculus? What are their requirements?
- How can network calculus computations be made more scalable? Where are fundamental limits for the network analysis? How do current software tools perform?
- What is the "killer" application for network calculus, and, in particular, for stochastic network calculus?
- How can network calculus' scope be extended to open up for new application domains?

Some (partial) answers to these important questions could be hinted at by the presentations and the subsequent discussions:

- Currently, some of the most promising application domains of (deterministic) network calculus were identified in industrial control, automotive and aerospace industries; also, interesting steps using (stochastic) network calculus in the modelling of smart energy grids were presented.
- The hardness of feedforward network analysis is by now understood, good heuristic approaches are on the way; however, cyclic dependencies and feedback systems are still open problems to some degree.
- The modelling of parallel systems using network calculus seems a promising building block to address novel attractive applications.
- Software tool support for network calculus, in particular for the stochastic version, is under construction and requires a community effort.

Looking over the fence: related methods. The research goals of network calculus and its methodologies, such as system performance evaluation, Markov chain analysis, or large deviations, intersect with those of other research communities. The objective of the session "Related Methods" was to create a forum where researchers from diverse research communities present their research approaches and discuss them with network calculus researchers. Thus, the session exposed the network calculus community to recent trends in system performance evaluation. Moreover, since speakers in this session had previously no or only limited exposure to network calculus, the session created an opportunity to disseminate the network calculus research agenda to other communities. The session was subtitled as "Looking over the fence", indicating an interest in learning new methodologies and the desire for cross- and interdisciplinary

interactions. The session featured speakers from four countries (Canada, France, Israel, USA), from three disciplines (mathematics, theoretical computer science, operations research), presenting recent research on approaches on topics such as robust queueing theory, adversarial queueing theory, and competitive analysis.

Also, the seminar comprised a one minute madness session for introduction and for statements on the network calculus, a breakout session for group work on promising future research topics in the network calculus, as well as a podium discussions on wireless network calculus. The discussions, viewpoints, and results that were obtained are also summarized in the full report.

We would like to thank all presenters, scribes, and participants for their contributions and lively discussions. Particular thanks go to the team of Schloss Dagstuhl for their excellent organization and support.

Fig. 7.6
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.19 Mixed Criticality on Multicore/Manycore Platforms

Organizers: Sanjoy K. Baruah, Liliana Cucu-Grosjean, Robert I. Davis, and Claire Maiza

Seminar No. 15121

Date: March 15–20, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.84

Participants: Yasmina Abdeddaim, Sebastian Altmeyer, James H. Anderson, Sanjoy K. Baruah, Marko Bertogna, Enrico Bini, Björn B. Brandenburg, David Broman, Alan Burns, Albert Cohen, Liliana Cucu-Grosjean, Robert I. Davis, Arvind Easwaran, Pontus Ekberg, Rolf Ernst, Sébastien Faucou, Nathan Fisher, Gerhard Fohler, Christopher D. Gill, Adriana Gogonel, Joel Goossens, Emmanuel Grolleau, Zhishan Guo, Pengcheng Huang, Leandro Soares Indrusiak, Kai Lampka, Björn Lisper, Claire Maiza, Alberto Marchetti-Spaccamela, Cristian Maxim, Dorin Maxim, Vincent Nelis, Roman Obermaisser, Gabriel Parmer, Sophie Quinton, Jan Reineke, Pascal Richard, Christine Rochange, Zoe Stephenson, Sebastian Stiller, Leen Stougie, Lothar Thiele, Suzanne van der Ster, Wang Yi

Real-time systems are characterised not only by the need for functional correctness, but also the need for timing correctness. Today, real-time embedded systems are found in many diverse application areas including; automotive electronics, avionics, and space systems. In these areas, technological progress is resulting in rapid increases in both software complexity and processing demands. To address the demand for increased processor performance, silicon vendors no longer concentrate on increasing processor clock speeds, as this approach has led to problems with high power consumption and excessive heat dissipation. Instead, technological development has shifted to multicore processors, with multiple CPUs integrated onto a single chip. The broad technology trend is towards much larger numbers of cores, referred to as manycore, requiring network-on-chip rather than bus interconnects.

Requirements on Size Weight and Power consumption, as well as unremitting cost pressures, are pushing developments in avionics and automotive electronics towards the adoption of powerful embedded multicore processors, with a longer term vision of migrating to manycore. With the adoption of such technology comes the opportunity to combine different applications on the same platform, potentially dramatically reducing assembly and production costs, while also improving reliability through a reduction in harnessing. Different applications may have different criticality levels (e.g. safety-critical, mission-critical, non-critical) designating the level of assurance needed against failure. For example, in automotive electronics, cruise control is a low criticality application, whereas electric steering assistance is of high criticality. In an aerospace context, flight control and surveillance applications in Unmanned Aerial Vehicles are of high and low criticality respectively. The very low acceptable failure rates (e.g. 10^{-9} failures per hour) for high criticality applications imply the need for significantly more rigorous and costly development and verification processes than required by low criticality applications.

Combining high and low criticality applications on the same hardware platform raises issues of time separation and composition; it must be possible to prevent the timing behaviour of high criticality applications from being disturbed by low criticality ones, otherwise both need to be engineered to the same rigorous and expensive standards. Simple methods of achieving this separation, such as time partitioning or allocation to different cores can however be wasteful of processing resources. They may require more expensive hardware than necessary, increasing production costs, which is something industry is strongly motivated to avoid. Time composability is needed so that the timing behaviour of applications, determined in isolation, remains valid when they are composed during system integration. Without time composability integration of complex applications would become infeasible expensive. The transformation of real-time embedded systems into mixed criticality multicore and manycore systems is recognised as a strategically important research area in Europe and the USA.

The seminar focused on the two key conflicting requirements of Mixed Criticality Systems: separation between criticality levels for assurance and sharing for resource efficiency, along with the related requirement of time composability. The key research questions addressed were:

- How to provide effective guarantees of real-time performance to applications of different criticality levels via intelligent sharing of resources while respecting the requirements for asymmetric separation / isolation between criticality levels?
- How to provide asymmetric time separation between applications with different levels of criticality so that the impact of lower criticality applications on those of higher criticality can be tightly bounded independent of the behaviour or mis-

behaviour of the former, without significantly compromising guaranteed real-time performance?
- How to provide time composability for applications of different criticality levels, so that the timing behaviour of applications determined in isolation remains valid when they are composed during system integration?

The sessions of the seminar were structured around a set of themes. Particular attention was given to the interfaces between themes, as these are the areas that can benefit most from improved understanding and collaboration. The discussion groups were organized around the following themes that correspond to research challenges in mixed criticality systems (MCS):
- Platforms and Experimental Evaluation;
- Worst-Case Execution Time;
- Criticality;
- Probabilistic.

Organization of the Seminar. The seminar took place from 15th to 20th March 2015. The first day started with a keynote talk by Prof. Alan Burns (University of York), one of the most influential researchers in the Real-Time Systems field over the last 25 years. Alan reviewed advances in MCS research and underlined current open problems. The first day ended with presentations and feedback on real implementations as well as identifying the main themes for group discussion.

The following three days started with presentations, which were followed by discussions either within the identified groups or in an open format.

The second day started with discussions about the motivation for mixed-criticality systems presented by three different participants. Different notations are used by different sub-communities and several presentations underlined these differences. An outline of the main ideas for probabilistic analysis of real-time systems provided the topics for the discussion group on probabilistic MCS.

The morning of the third day commenced with discussions on the relation between time and MCS, which continued into the afternoon's hiking activity.

Starting from the fourth day a slot dedicated to anonymous mixed criticality supporters was added to the program allowing researchers new to the topic to identify open problems in MCS from the perspective of their different domains.

As detailed later in this report, the seminar enabled the real-time community to make important progress in articulating and reaching a common understanding on the key open problems in mixed criticality systems, as well as attracting new researchers to these open problems. The seminar also provided an ideal venue for commencing new collaborations, a number of which are progressing towards new research publications.

The seminar has helped define a research agenda for the coming years that could be supported by follow-up events, given the strong interest expressed by the participants of this seminar.

As organizers, we would like to thank Prof. Reinhard Wilhelm for encouraging us to submit the seminar proposal, Dagstuhl's Scientific Directorate for allowing us to run a seminar on mixed criticality systems, and to the staff at Schloss Dagstuhl for their superb support during the seminar itself. Finally, we would like to thank all of the participants for their strong interaction, presentations, group discussions, and work on open problems, sometimes into the early hours of the morning. We were very pleased to hear about the progress of new found collaborations, and to receive such positive feedback about the seminar itself. Thank you to everyone who participated for a most enjoyable and fruitful seminar.

7.20 Formal Models of Graph Transformation in Natural Language Processing

Organizers: Frank Drewes, Kevin Knight, and Marco Kuhlmann

Seminar No. 15122

Date: March 15–20, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.143

Participants: Daniel Bauer, Suna Bensch, Henrik Björklund, Johanna Björklund, David Chiang, Frank Drewes, Petter Ericson, Daniel Gildea, Karl Moritz Hermann, Berthold Hoffmann, Peter Jonsson, Laura Kallmeyer, Kevin Knight, Alexander Koller, Marco Kuhlmann, Adam Lopez, Andreas Maletti, Jonathan May, Mark Minas, Joakim Nivre, Stephan Oepen, Detlef Plump, Giorgio Satta, Natalie Schluter, Mark Steedman, Christoph Teichmann, Brink van der Merwe, Heiko Vogler, Daniel Zeman

Strings are fundamental data structures in natural language processing (NLP). Weighted finite-state string acceptors and transducers, first introduced as theoretical constructs, have proven their worth in speech recognition, part-of-speech tagging, transliteration, and many other applications. The string automaton framework provides efficient generic algorithms for composition, bidirectional application, k-best extraction, determinization, minimization, parameter tuning, etc. These algorithms have been packaged in software toolkits that form the core of many state-of-the-art systems.

Tree automata go further in permitting large-scale, syntactically-motivated re-ordering of subtrees. They were originally devised to help formalize Chomsky's linguistic theories, but their subsequent development was largely disconnected from NLP practice. In 2005, tree automata theorists and machine translation (MT) practitioners began working together to come up with a new kind of statistical MT system based on tree automata. This led to some of the best practical results in common evaluations of MT quality, and syntactic methods are now used in industrial MT systems. This work at the intersection of tree automata and NLP created vibrant new research directions for both areas.

Nowadays, *graphs* are becoming an even more general fundamental data structure in practical NLP. Classic feature structures can be seen as rooted, directed, edge- and leaf-labeled graphs. Recent work in dependency parsing produces graphs rather than trees. New work in deep semantic annotation organizes logical meanings into directed graph structures, and several efforts are now being made that in the near future will yield large amounts of linguistic data annotated with these representations. Formal models of *graph transformation* are therefore of fundamental importance for the development of practical systems for these tasks. The situation is familiar: there exists a formal theory of graph transformation, but this theory is largely disconnected from research and practice in NLP.

The theory of graph transformation studies rule-based mechanisms for the manipulation of graphs. A particularly well-studied subject within the area of graph transformation, and one that has received quite some attention recently within the NLP community, are *context-free graph grammars*. These grammars have many nice properties in common with context-free phrase structure grammars, but are considerably more powerful and versatile; in particular, they can be used to generate context-sensitive string languages (when strings are represented as chain graphs). The price of this expressiveness is a higher computational complexity; in particular, there are context-free graph languages for which parsing is NP-complete. This has triggered research on specialized, more efficient algorithms for restricted classes of graphs. A well-known result in this area is that many in general intractable problems on graphs become solvable in polynomial time when restricted to graphs of bounded tree-width.

With the number of interesting applications and the amount of available data quickly increasing, there is a clear need for the NLP community to acquire knowledge about formal models of graph processing, as such models can greatly simplify practical systems, by providing a uniform knowledge representation and efficient, generic algorithms for inference. Unfortunately, most NLP researchers are unaware of the rich literature on graph transformation, and even those who are find it hard to connect it to their own work. Conversely, few researchers in graph transformation are aware of the new applications of their research within natural language processing, the characteristic properties of the available data, the specific desiderata of these applications, and the research problems that are posed by them.

The overall goal of the seminar was to bring the various research communities together to assess the state of the art, identify areas of common interest, and pave the way for future collaborations. We think that this goal was reached to a very high degree, which will be a major factor in the creation of a new interdisciplinary research community.

Fig. 7.7
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.21 Normative Multi-Agent Systems

Organizers: Amit K. Chopra, Leon van der Torre, Harko Verhagen, and Serena Villata

Seminar No. 15131

Date: March 22–27, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.3.162

Participants: Huib Aldewereld, Diego Agustin Ambrossio, Matteo Baldoni, Simon Caton, Amit K. Chopra, Rob Christiaanse, Silvano Colombo Tosatto, Célia da Costa Pereira, Christoph Dorn, Hein Duijf, Corinna Elsenbroich, Dov M. Gabbay, Aditya K. Ghose, Guido Governatori, Joris Hulstijn, Llio Humphreys, Franziska Klügl, Ho-Pun Lam, Beishui Liao, Daniel Moldt, Robert Muthuri, Luis Gustavo Nardin, Martin Neumann, Pablo Noriega, Julian Padget, Adrian Paschke, Gabriella Pigozzi, Axel Polleres, Livio Robaldo, Victor Rodriguez Doncel, Antonino Rotolo, Ken Satoh, Judith Simon, Munindar P. Singh, Xin Sun, Viviane Torres da Silva, Leon van der Torre, Wamberto Vasconcelos, Harko Verhagen, Serena Villata

The multi-disciplinary workshop on Normative Multi Agents attracted leading international scholars from different research fields (e.g., theoretical computer science, programming languages, cognitive sciences and social sciences). The workshop was organized as follows: the organizers identified three relevant themes of research covering a wide and comprehensive spectrum of topics in the filed of Normative Agents, namely Social Computing, Governance, and Agreement Technologies. In the months preceding the workshop the chairs collected material from the participants. During the first day each participant present herself to the audience, and the chairs introduced the goal of the seminar, i.e., writing an handbook of Normative Multiagent Systems based on the roadmap produced during the previous edition of the Seminar, and the discussions during the current one. The participants were divided in groups corresponding to the areas identified as relevant in the field of Normative Multiagent Systems. Four invited talks have been proposed by scholars from different areas in the field, targeting in particular the three themes of the Seminar and an overview about Normative Multiagent Systems. The format was well received by the participants and conducive to discussion. It gave them the opportunity to give very focused presentations while keeping the audience attention. During the morning sessions, we started with an invited talk and we continued with short presentations by the Seminar participants about their personal contribution to Normative Multi-Agents (plus some time for QA). The afternoon sessions, other the contrary, were dedicated to group work and group discussions. The aim of these sessions was to build consensus material of the specific topics and to identify fundamental research directions. The material is expected to be refined and to be articulated in chapters intended as a first step for the development for the handbook for this emerging area of computer-science with close interactions with other disciplines.

Results

During the seminar, participants split in different working groups, centered around discussion themes relevant to NorMAS. Each working group was further divided into smaller working groups, each of which worked on specific topics. In the following paragraphs there is a summary of the discussion held by each working group.

Logic and reasoning. This theme included subgroups on topics such as *deontic logic*, *argumentation*, *computation approaches*, *motivational attitudes*, *social games*, and *emotions*.

Modeling. This theme included subgroups on issues such as *taxonomies*, *law*, *conflicts*, and *norm dynamics*.

Engineering. This theme included subgroups on themes such as *interactions*, *agent programming*, *agent architecture*, *data-driven norms*, *institutions and technology*, and *reference architectures*.

Simulation. This theme discussed issues of simulating multiagent systems to understand norm dynamics such as *emergence* and *diffusion*.

Applications. This theme included subgroups on killer applications for norms. Identified applications included *governance*, *audit control*, *cybersecurity*, *jurisinformatics*, and *sociotechnical systems*.

Each subgroup presented its findings twice to the entire seminar. Each subgroup identified past work, connections to other subgroups, and future work. Based on their presentations, we decided that each subgroup should write a chapter on its topic. This chapter will become part of a Handbook of Normative Multiagent Systems. This is in line with the roadmap produced during the previous edition of the Seminar and the discussions

held during the present Seminar. The handbook will be an authoritative and detailed introduction for anyone seeking information on normative multiagent systems. The handbook will give a historical overview, present a survey of established techniques and open challenges, and discuss applications and directions. Our aim is to have to handbook sent for publication in a year's time. We already have a publisher lined up (College Publications).

7.22 Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations

Organizers: David Hutchison, Klara Nahrstedt, Marcus Schöller, Indra Spiecker gen. Döhmann, and Markus Tauber

Seminar No. 15151

Date: April 7–10, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.1

Participants: Ali Alshawish, Silvia Balaban, Saleem Bhatti, Roland Bless, Marcus Brunner, György Dan, Jerker Delsing, Thilo Ewald, Andreas Fischer, David Hutchison, Youki Kadobayashi, Graham Kirby, Helmut Leopold, Andreas Mauthe, Simon Oechsner, Frank Pallas, Oliver Raabe, Alberto Egon Schaeffer-Filho, Burkhard Schafer, Marcus Schöller, Sean W. Smith, Christoph Sorge, Indra Spiecker gen. Döhmann, James P. G. Sterbenz, Burkhard Stiller, Markus Tauber, Gene Tsudik, Pal Varga, Edgar R. Weippl

This report documents the programme and the outcomes of Dagstuhl Seminar 15151 on "Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations". The main objective of the Seminar was to bring together researchers from different disciplines in order to establish a research agenda for securing services-to-come in our increasingly connected world. The backgrounds and interests of the participants included i) techno-legal, ii) resilience and systems security, and iii) socio-technical topics. The use case domains that were discussed covered the Internet of Things (IoT) as well as Cloud-based applications in which flexible service composition is paramount. We started the seminar using four introductory talks covering respectively the "big picture", the legal viewpoint, the technical viewpoint, and the organisational viewpoint. From this beginning, we derived initial research questions in small groups, and these questions and issues arising were then consolidated and refined into the resulting material that is presented below.

The opening speakers were the following:

- Helmut Leopold, Head of the Digital Safety and Security Department at the Austrian Institute of Technology, who presented the "big picture", i.e. where our connected world is heading;
- Burkhard Schafer, Professor of Computational Legal Theory at the University of Edinburgh, who presented his viewpoint on legal challenges within our ever interconnected society;
- Thilo Ewald from Microsoft Deutschland GmbH, who explained his viewpoint on the organisational challenges in today's world;
- Marcus Brunner, Head of Standardization in the strategy and innovation department of Swisscom, presented his viewpoint on technological developments in designing and building flexible networked systems.

From this starting point we derived initial research questions in small groups. The organising team reviewed intermediate results and re-balanced groups and most significantly identified the core questions to work on. The groups were between 4 and 6 people at any time, and a good balance was maintained across the representatives of legal, organisational and technological experts and between the groups. The resulting questions and issues were:

1. How to enable Resilience, by design, of composable flexible systems [1]?
2. What is the role of law in supporting resilience, privacy [2] and security?
3. Traceability of (personal and non-personal) data in service provision?
4. How can we improve the perception of assurance [3], privacy, security and resilience for the end-user?
5. What constitutes a security problem?
6. How to deal with unforeseen new context of usage?

These questions were crucial, in that they formed the basis for the bulk of group discussions throughout the second and third days of the Seminar. Therefore, the organisers took great care – and a great deal of time during the first evening – formulating these questions, together with the related issues. At the start of the second day, these questions and issues were presented to the groups, who were invited to comment on them. The groups were invited to add their own interpretation, and to identify additional issues during their discussions. During the subsequent periods – broken up by refreshments and lunch – the organisers checked that the groups appeared to be productive and harmonious (which on both counts they turned out to be). Each group was asked to record the essence of their discussions, and conclusions, and to pass these to the organisers by the end of the Seminar. Every group did some additional work after the Seminar, and the report assembled here reflects the hard work of the participants as well as the organisers, during the Seminar itself and in the days that followed.

References

1 James P. G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Comput. Netw.*, 54(8):1245–1265, June 2010.

2 Burkhard Schafer. All changed, changed utterly? *Datenschutz und Datensicherheit – DuD*, 35(9):634–638, 2011.

3 Aleksandar Hudic, Markus Tauber, Thomas Lorunser, Maria Krotsiani, George Spanoudakis, Andreas Mauthe, and Edgar R. Weippl. A multi-layer and multitenant cloud assurance evaluation methodology. In *Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on*, pages 386–393. IEEE, 2014.

7.23 Machine Learning with Interdependent and Non-Identically Distributed Data

Organizers: Trevor Darrell, Marius Kloft, Massimiliano Pontil, Gunnar Rätsch, and Erik Rodner

Seminar No. 15152

Date: April 7–10, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.18

Participants: Shai Ben-David, Gilles Blanchard, Trevor Darrell, Joachim Denzler, Philipp Drewe, Mario Fritz, Judy Hoffman, Josef Kittler, Marius Kloft, Brian Kulis, Christoph H. Lampert, Soeren Laue, Alessandro Lazaric, Victor Lempitsky, Christoph Lippert, Stephan Mandt, Shin Nakajima, Francesco Orabona, Massimiliano Pontil, Gunnar Rätsch, Erik Rodner, Kate Saenko, Tobias Scheffer, Dino Sejdinovic, Fei Sha, Oliver Stegle, Ingo Steinwart, Ilya Tolstikhin, Ruth Urner, Alexander Zimin

The seminar broadly dealt with *machine learning*, the area of computer science that concerns developing computational methods using data to make accurate predictions. The classical machine learning theory is built upon the assumption of independent and identically distributed random variables. In practical applications, however, this assumption is often violated, for instance, when training and test data come from different distributions (dataset bias or domain shift) or when the data exhibits temporal or spatial correlations. In general, there are three major reasons why the assumption of independent and identically distributed data can be violated:

1. The draw of a data point influences the outcome of a subsequent draw (inter-dependencies).
2. The distribution changes at some point (non-stationarity).
3. The data is not generated by a distribution at all (adversarial).

The seminar focused on the scenarios (a) and (b). This general research direction comprises several subfields of machine learning: transfer and multi-task learning, learning with interdependent data, and two application fields, that is, visual recognition and computational biology. Both application areas are not only two of the main application areas for machine learning algorithms in general, but their recognition tasks are often characterized by multiple related learning problems that require transfer and multitask learning approaches. For example, in visual recognition tasks, object categories are often visually related or hierarchically organized, and tasks in computational biology are often characterized by different but related organisms and phenotypes. The problems and techniques discussed during the seminar are also important for other more general application areas, such as scientific data analysis or data-oriented decision making.

Results of the Seminar and Topics Discussed

In the following, the important research fields related to the seminar topic are introduced and we also give a short list of corresponding research questions discussed at the seminar. In contrast to other workshops and seminars often associated with larger conferences, the aim of the Dagstuhl seminar was to reflect on open issues in each of the individual research areas.

Foundations of Transfer Learning. Transfer Learning (TL) [2, 18] refers to the problem of retaining and applying the knowledge available for one or more source tasks, in order to efficiently develop an hypothesis for a new target task. Each task may contain common (domain adaptation [10, 25]) or different label sets (across category transfer). Most of the effort has been devoted to binary classification [23], while interesting practical transfer problems are often intrinsically multi-class and the number of classes can increase in time [17, 22]. Accordingly the following research questions arise:

- How to formalize knowledge transfer across multi-class tasks and provide theoretical guarantees on this setting?
- Moreover, can inter-class transfer and incremental class learning be properly integrated?
- Can learning guarantees be provided when the adaptation relies only on pre-trained source hypotheses without explicit access to the source samples, as it is often the case in real world scenarios?

Foundations of Multi-task Learning. Learning over multiple related tasks can outperform learning each task in isolation. This is the principal assertion of Multi-task learning (MTL) [1, 3, 7] and implies that the learning process may benefit from common information shared across the tasks. In the simplest case, the transfer process is symmetric and all the tasks are

considered as equally related and appropriate for joint training. Open questions in this area are:

- What happens when the condition of equally related tasks does not hold, e.g., how to avoid negative transfer?
- Moreover, can non-parametric statistics [27] be adequately integrated into the learning process to estimate and compare the distributions underlying the multiple tasks in order to learn the task similarity measure?
- Can recent semi-automatic methods, like deep learning [9] or multiple kernel learning [4, 11–13], help to get a step closer towards the complete automatization of multi-task learning, e.g., by learning the task similarity measure?
- How can insights and views of researcher be shared across domains (e.g., regarding the notation of *source task selection* in reinforcement learning)?

Foundations of Learning with Inter-dependent Data. Dependent data arises whenever there are inherent correlations in between observations. For example, this is to be expected for time series, where we would intuitively expect that instances with similar time stamps have stronger dependencies than ones that are far away in time. Another domain where dependent data occurs are spatially-indexed sequences, such as windows taken from DNA sequences. Most of the body of work on machine learning theory is on learning with i.i.d. data. Even the few analyses (e.g., [28]) allowing for "slight" violations of the assumption (mixing processes) analyze the same algorithms as in the i.i.d. case, while it should be clear that also novel algorithms are needed to most effectively adapt to rich dependency structures in the data. The following aspects have been discussed during the seminar:

- Can we develop algorithms that exploit rich dependency structures in the data?
- Do such algorithms enjoy theoretical generalization guarantees?
- Can such algorithms be phrased in a general framework in order to jointly analyze them?
- How can we appropriately measure the degree of inter-dependencies (theoretically) such that it can be also empirically estimated from data (overcoming the so-called *mixing* assumption)?
- Can theoretical bounds be obtained for more practical dependency measures than mixing?

Visual Transfer and Adaptation. Visual recognition tasks are one of the main applications for knowledge transfer and adaptation techniques. For instance, transfer learning can put to good use in the presence of visual categories with only a few number of labels, while across category transfer can help to exploit training data available for related categories to improve the recognition performance [14, 20–22]. Multi-task learning can be applied for learning multiple object detectors [30] or binary image classifiers [19] jointly, which is beneficial because visual features can be shared among categories and tasks. Another important topic is domain adaptation, which is very effective in object recognition applications [24], where the image distribution used for training (source domain) is different from the image distribution encountered during testing (target domain). This distribution shift is typically caused by a data collection bias. Sophisticated methods are needed as in general the visual domains can differ in a combination of (often unknown) factors including scene, object location and pose, viewing angle, resolution, motion blur, scene illumination, background clutter, camera characteristics, etc. Recent studies have demonstrated a significant degradation in the performance of state-of-the-art image classifiers due to domain shift from pose changes [8], a shift from commercial to consumer video [5, 6, 10], and, more generally, training datasets biased by the way in which they were collected [29].

The following open questions have been discussed during the seminar:

- Which types of representations are suitable for transfer learning?
- How can we extend and update representations to avoid negative transfer?
- Are current adaptation and transfer learning methods efficient enough to allow for large-scale continuous visual learning and recognition?
- How can we exploit huge amounts of unlabeled data with certain dependencies to minimize supervision during learning and adaptation?
- Are deep learning methods already compensating for common domain changes in visual recognition applications?

Application Scenarios in Computational Biology. Non-i.i.d. data arises in biology, e.g., when transferring information from one organism to another or when learning from multiple organisms simultaneously [31]. A scenario where dependent data occurs is when extracting local features from genomic DNA by running a sliding window over a DNA sequence, which is a common approach to detect transcription start sites (TSS) [26]. Windows close by on the DNA strand – or even overlapping – show stronger dependencies than those far away. Another application scenario comes from statistical genetics. Many efforts in recent years focused on models to correct for population structure [16], which can arise from inter dependencies in the population under investigation. Correcting for such rich dependency structures is also a challenge in prediction problems in machine learning [15]. The seminar brought ideas together from the different fields of machine learning, statistical genetics, Bayesian probabilistic modeling, and frequentist statistics. In particular, we discussed the following open research questions:

- How can we empirically measure the degree of inter-dependencies, e.g., from a kinship matrix of patients?
- Do theoretical guarantees of algorithms (see above) break down for realistic values of "the degree of dependency"?
- What are effective prediction and learning algorithms correcting for population structure and inter-dependencies in general and can they be phrased in a general framework?
- What are adequate benchmarks to evaluate learning with non-i.i.d. data?
- How can information be transferred between organisms, taking into account the varying noise level and experimental conditions from which data are derived?
- How can non-stationarity be exploited in biological applications?
- What are promising applications of non-i.i.d. learning in the domains of bioinformatics and personalized medicine?

Conclusion. The idea of the seminar bringing together people from theory, algorithms, computer vision, and computational biology, was very successful, since many discussions and joint research questions came up that have not been anticipated in the beginning. These aspects were not completely limited to non-i.i.d. learning and also touched ubiquitous topics like learning with deeper architectures. It was the agreement of all participants that the seminar should be the beginning of an ongoing series of longer Dagstuhl seminars focused on non-i.i.d. learning.

References

1 A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.

2 J. Baxter. A model of inductive bias learning. *J. of Artificial Intelligence Research*, 12:149–198, 2000.

3 R. Caruana. Multitask learning. *Machine Learning*, 28:41–75, July 1997.

4 C. Cortes, M. Kloft, and M. Mohri. Learning kernels using local rademacher complexity. In *Advances in Neural Information Processing Systems 26*, 2013.

5 L. Duan, I. W. Tsang, D. Xu, and S. J. Maybank. Domain transfer svm for video concept detection. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2009.

6 L. Duan, D. Xu, I. Tsang, and J. Luo. Visual event recognition in videos by learning from web data. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2010.

7 T. Evgeniou and M. Pontil. Regularized multi–task learning. In *Proc. of the 10 ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining*, pp. 109–117. ACM, 2004.

8 A. Farhadi and M. Kamali Tabrizi. Learning to recognize activities from the wrong view point. In *Proc. of the European Conf. on Computer Vision*, 2008.

9 G. E. Hinton, S. Osindero, and Y.-W. Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7):1527–1554, July 2006.

10 J. Hoffman, E. Rodner, J. Donahue, T. Darrell, and K. Saenko. Efficient learning of domain-invariant image representations. In *Int'l Conf. on Learning Representations (ICLR)*, 2013.

11 M. Kloft and G. Blanchard. On the convergence rate of ℓ_p-norm multiple kernel learning. *Journal of Machine Learning Research*, 13:2465–2502, Aug 2012.

12 M. Kloft, U. Brefeld, S. Sonnenburg, and A. Zien. Lp-norm multiple kernel learning. *Journal of Machine Learning Research*, 12:953–997, Mar 2011.

13 G. Lanckriet, N. Cristianini, L. E. Ghaoui, P. Bartlett, and M. I. Jordan. Learning the kernel matrix with semi-definite programming. *Journal of Machine Learning Research*, 5:27–72, 2004.

14 F.-F. Li, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):594–611, 2006.

15 L. Li, B. Rakitsch, and K. M. Borgwardt. ccsvm: correcting support vector machines for confounding factors in biological data classification. *Bioinformatics [ISMB/ECCB]*, 27(13):342–348, 2011.

16 C. Lippert, J. Listgarten, Y. Liu, C. M. Kadie, R. I. Davidson, and D. Heckerman. FaST linear mixed models for genome-wide association studies. *Nat Meth*, 8(10):833–835, October 2011.

17 J. Luo, T. Tommasi, and B. Caputo. Multiclass transfer learning from unconstrained priors. In *Proc. of the Int'l Conf. on Computer Vision (ICCV)*, pp. 1863–1870, 2011.

18 S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.

19 A. Quattoni, M. Collins, and T. Darrell. Transfer learning for image classification with sparse prototype representations. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 1–8, 2008.

20 E. Rodner and J. Denzler. Learning with few examples by transferring feature relevance. In *Proc. of the 31st Annual Symp. of the German Association for Pattern Recognition*, pp. 252–261, 2009.

21 E. Rodner and J. Denzler. One-shot learning of object categories using dependent gaussian processes. In *Proc. of the 32nd Annual Symp. of the German Association for Pattern Recognition*, pp. 232–241, 2010.

22 E. Rodner and J. Denzler. Learning with few examples for binary and multiclass classification using regularization of randomized trees. *Pattern Recognition Letters*, 32(2):244–251, 2011.

23 U. Rückert and M. Kloft. Transfer learning with adaptive regularizers. In *ECML/PKDD (3)*, pp. 65–80, 2011.

24 K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In *European Conf. on Computer Vision*, pp. 213–226, 2010.

25 G. Schweikert, C. Widmer, B. Schölkopf, and G. Rätsch. An empirical analysis of domain adaptation algorithms for genomic sequence analysis. In *Advances in Neural Information Processing Systems 21*, pp. 433–1440, 2009.

26 S. Sonnenburg, A. Zien, and G. Rätsch. Arts: Accurate recognition of transcription starts in human. *Bioinformatics*, 22(14):e472–e480, 2006.

27 B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research*, 11:1517–1561, 2010.

28 I. Steinwart, D. R. Hush, and C. Scovel. Learning from dependent observations. *J. Multivariate Analysis*, 100(1):175–194, 2009.

29 A. Torralba and A. Efros. Unbiased look at dataset bias. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, 2011.

30 A. Torralba, K. P. Murphy, and W. T. Freeman. Sharing features: efficient boosting procedures for multiclass object detection. In *IEEE Conf. on Computer Vision and Pattern Recognition Vol. 2*, pp. 762–769, 2004.

31 C. Widmer, M. Kloft, and G. Rätsch. Multi-task learning for computational biology: Overview and outlook. In *B. Schoelkopf, Z. Luo, and V. Vovk (eds.), Empirical Inference – Festschrift in Honor of Vladimir N. Vapnik*, 2013.

Fig. 7.8
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.24 Advanced Stencil-Code Engineering

Organizers: Christian Lengauer, Matthias Bolten, Robert D. Falgout, and Olaf Schenk

Seminar No. 15161

Date: April 12–17, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.56

Participants: Sven Apel, Mauro Bianco, Matthias Bolten, Uday Bondhugula, Rezaul Chowdhury, Simplice Donfack, Christian Engwer, Robert D. Falgout, Franz Franchetti, Michael Freitag, Stephanie Friedhoff, Francisco Jose Gaspar, Björn Gmeiner, Alexander Grebhahn, Armin Größlinger, Tobias Grosser, Gundolf Haase, Georg Hager, Frank Hannig, Juraj Kardos, Paul H. J. Kelly, Hans-Peter Kersken, Harald Köstler, Stefan Kronawitter, Sebastian Kuckuk, Bradley C. Kuszmaul, Christian Lengauer, Dmitry Mikushin, Marcus Mohr, Carlos Osuna Escamilla, Ekanathan Palamadai Natarajan, Dirk Pflüger, Louis-Noël Pouchet, Hannah Rittich, Carmen Rodrigo Cardiel, Ulrich Rüde, P. Saday Sadayappan, Olaf Schenk, Christian Schmitt, Armando Solar-Lezama, Jürgen Teich, Wim Vanroose, Christian Waluga, Gabriel Wittum, Barbara Wohlmuth, Ulrike Meier Yang

Stencil Codes

Stencil codes are compute-intensive algorithms, in which data points arranged in a large grid are being recomputed repeatedly from the values of data points in a predefined neighborhood. This fixed neighborhood pattern is called a stencil. Stencil codes see wide-spread use in computing the discrete solutions of partial differential equations and systems composed of such equations. Connected to the implementation of stencil codes is the use of efficient solver technology, i.e., iterative solvers that rely on the application of a stencil and that provide good convergence properties like multigrid methods. Major application areas are the natural sciences and engineering. Although, in many of these applications, unstructured adaptive discretizations are employed for an efficient use of exascale supercomputers whose architectures possibly include accelerators or are of a heterogeneous nature, the use of structured discretizations and, thus, stencil codes has turned out to be helpful.

Stencil codes come in large varieties: there are many thousands! Deriving each of them individually, even if by code modification from one another, is not practical. The goal of the seminar is to raise the level of abstraction for application programmers significantly and to support this raise with an automated software technology that generates highly efficient massively parallel implementations which are tuned to the specific problem at hand and the execution platform used.

Research Challenges

Stencil codes are algorithms with a pleasantly high regularity: the data structures are higher-dimensional grids and the computations follow a static, locally contained dependence pattern and are typically arranged in nested loops with linearly affine bounds. This invites massive parallelism and raises the hope for easily achieved high performance. However, serious challenges remain:

- Because of the large numbers and varieties of stencil code implementations, deriving each of them individually, even if by code modification from one another, is not practical. Not even the use of program libraries is practical; instead, a domain-specific metaprogramming approach is needed.
- Reaching petascale to exascale execution speed is a challenge in the frequently used so-called multigrid algorithms, which work on a hierarchy of increasingly larger grids. The coarse grids in the upper part of the hierarchy are too small for massive parallelism.
- Efficiency, i.e., a high ratio of speedup to the degree of parallelism, is impaired by the low mathematical density, i.e., the low ratio of computation steps to data transfers of stencil codes.
- An inappropriate use of the execution platform may act as a performance brake.

Stencil-code engineering has received increased attention in the last few years, which is evidenced by the appearance of a number of stencil-code programming languages and frameworks. To reach the highest possible execution speed and to conserve hardware resources and energy, the stencil code must be tuned cleverly to the specific application problem at hand and the execution platform used. One approach that could be followed has been demonstrated by the previous U.S. project SPIRAL, whose target was the domain of linear transforms: domain-specific optimization at several levels of abstraction – from the mathematical equations over an abstract, domain-specific program and, in further steps, to the actual target code on the execution platform used. At each level, one makes aggressive use of knowledge of the problem and platform and employs up-to-date, automated software technology suitable for that level.

Questions and Issues Addressed

The charter of the seminar was to foster international cooperation in the development of a radically new, automatic, optimizing software technology for the effective and flexible exploitation of massively parallel architectures for dedicated, well delineated problem domains.

The central approaches in achieving this technology are:

- the aggressive use of domain knowledge for optimization at different levels of abstraction
- the exploitation of commonalities and variabilities in application codes via product-line technology and domain engineering
- the use of powerful models for program optimization, like the polyhedron model for loop parallelization and feature-orientation for software product lines

The application domain investigated in the seminar was stencil codes. It is envisaged that the approach can be ported to other well delineated domains – of course, with the substitution of suitable domain-specific content.

Among the issues discussed were:

- What are suitable abstraction, modularization, composition and generation mechanisms for stencil codes?
- What are the appropriate language features of a domain-specific language for stencil codes?
- What are the commonalities and variabilities of stencil codes?
- What are the computational performance barriers, especially, of multigrid methods using stencils and how can they be overcome?
- What are the performance barriers caused by data exchanges and how can they be overcome? How can communication be avoided in multilevel algorithms?
- What are the roles of nested loops and divide-and-conquer recursions in stencil codes?
- How can other solvers and preconditioners benefit from autotuned stencil codes?
- What role should techniques like autotuning and machine learning play in the optimization of stencil codes?
- What options of mapping stencil codes to a heterogeneous execution platform exist and how can an educated choice be made?
- Which techniques can be employed to make clever use of large-scale hybrid architectures, e.g., by the combination of multigrid with mathematical domain decomposition?

On the informatics side, one important role of the seminar was to inform the international stencils community about the techniques used in ExaStencils: software product lines, polyhedral loop optimization and architectural metaprogramming. Equally important was for ExaStencils members to learn about the experiences made with other techniques like divide-and-conquer, multicore optimization in parallel algorithms or autotuning. The application experts contributed to a realistic grounding of the research questions.

On the mathematics side, the seminar fostered the cooperation of experts in parallel solver technology with the groups from informatics to enable them to make use of the advanced techniques available. Further, different strategies for improving the scalability of iterative methods were discussed and the awareness of the opportunities and complexities of modern architectures in the numerical mathematics community was advanced.

7.25 Software and Systems Traceability for Safety-Critical Projects

Organizers: Jane Cleland-Huang, Sanjai Rayadurgam, Patrick Mäder, and Wilhelm Schäfer

Seminar No. 15162

Date: April 12–17, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.76

Participants: Markus Borg, Jane Cleland-Huang, Krzysztof Czarnecki, Christopher Gerking, Paul Grünbacher, Lars Grunske, Kai Höfig, Patrick Mäder, Shiva Nejati, Leon J. Osterweil, Mona Rahimi, Sanjai Rayadurgam, Gilbert Regan, Patrick Rempel, Mehrdad Sabetzadeh, Nicolas Sannier, Wilhelm Schäfer, Andres Toom, Marc Zeller

Safety-critical systems, defined as systems whose "failure could result in loss of life, significant property damage, or damage to the environment"[40], pervade our society. Developing software is a challenging process. Not only must the software deliver the required features, but it must do so in a way that ensures that the system is safe and secure for its intended use. To this end safety-critical systems must meet stringent guidelines before they can be approved or certified for use. For example, software developed for the aerospace industry must comply to the ISO12207 and/or the DO-178B/C guidelines, while software developed for European railway communication, signaling, and processing systems, must comply to EN50128. Most guidelines prescribe a set of steps and deliverable documents that focus around planning, analysis and design, implementation, verification and validation, configuration management, and quality assurance activities. In addition they often provide specific guidelines for the creation and use of traceability in the project. For example, depending upon the criticality level of a requirement, the US Federal Aviation Authority guideline DO-178B requires traceability from requirements to design, and from requirements to source code and executable object code.

In practice, traceability is achieved through the creation and use of *trace links*, defined by the Center of Excellence for Software and Systems Traceability[41] as "specified associations between pair of artifacts, one comprising the source artifact and one comprising the target artifact". Software traceability serves an important role in demonstrating that a delivered software system satisfies its software design constraints and mitigates all identified hazards. When correct, traceability demonstrates that a rigorous software development process has been established and systematically followed. Current guidelines, in many safety-critical industries, prescribe traceability for two reasons. First, as an indirect measure that good practice has been followed, the general idea being that traceability information serves as an indicator that design and production practices were conducted in a sound fashion; and second, as a more direct measure, to show that specific hazards have been explored, potential failure modes identified, and that the system is designed and implemented in a "demonstrably rational way".

Unfortunately, there is a significant gap between prescribed and actual traceability. An analysis of the traceability information submitted by various organizations to the US Food and Drug Administration (FDA) as part of the medical device approval process [42], showed a significant *traceability gap* between the traceability expectations as laid out in the FDA's "Guidance for the Content of Premarket Submissions for Software Contained in Medical Devices", and the traceability data documented in the submissions. While all of the submissions made some attempt to satisfy the FDA's traceability guidelines, serious deficiencies were found in almost all the submissions in terms of missing traceability paths, missing and redundant links, and problems in trace granularity. These deficiencies made it very difficult to understand the rationale for individual links. A more recent systematic analysis of seven software projects that originated from four different domains (automotive, aviation, medical, and space)

[40] *Failure Analysis and the Safety-Case Lifecycle*, W. S. Greenwell, E. A. Strunk, and J. C. Knight in Human Error, Safety and Systems Development, 2004, http://dx.doi.org/10.1007/1-4020-8153-7_11.
[41] Center of Excellence for Software and Systems Traceability (http://www.CoEST.org)
[42] *Strategic traceability for safety-critical projects*, P. Mäder, P. L. Jones, Y. Zhang, and J. Cleland-Huang, IEEE Software, 30(3):58–66, http://dx.doi.org/10.1109/MS.2013.60.

revealed similar problems [43]. The provided software development artifacts were analyzed with respect to four technical guideline documents (ISO 26262-6, DO-178B, FDA Guide for Submissions, ECSS-E-40), where each document is a representative guideline of one of the four domains.

Problems are exacerbated in the systems engineering domain in which core concepts and designs are often documented across multiple models, each of which might depict a single viewpoint or perspective of the system. For example, the system might include separate models for functional and behavioral requirements, software components, electrical components, thermodynamics, and mechanical components. Furthermore, although each of these perspectives is modeled separately in isolation from one another, they interact to produce the final behavior of the system. Traceability solutions must extend across these heterogeneous models. Deficiencies in traceability are certainly not new. As far back as 1995, Gotel et al. identified several different traceability problems and attributed them to poor coordination, lack of perceived benefits, time to market pressures, and lack of sufficient tooling. These problems observed almost 20 years ago, continue to plague the traceability landscape today, meaning that the *traceability gap* between what is prescribed and what is practiced is still very real.

Given that the software and systems engineering communities have been unable to solve this problem in over 20 years, it seems prudent to reexamine traceability needs and their prescribed solutions. Within this Dagstuhl seminar, we engaged software and systems engineering researchers and practitioners from the safety-critical domain alongside traceability experts, in highly focused discussions. The aim was to gain a deeper understanding of exactly what traceability is needed for safety-critical systems, and to identify practical and achievable solutions. To the best of our knowledge this was the first time researchers from the safety-critical and traceability domains came together in a dedicated forum to tackle this problem.

We started the week with a number of more general presentations and discussions from experts in the respective areas to form a common understanding for later discussions. Subsequently, the seminar continued with shorter talks focusing on a variety of specific aspects of open challenges and potential solutions accompanied by intensive and highly interactive discussions. In parallel, we parted for about one third of the time into four focus groups working on what had been identified as the most relevant and urgent challenges for closing the traceability gap. The four areas of focus were: tracing qualities, traceability in the context of models and tools, cost-benefit and stakeholder perspectives, and traceability in the context of evolution and change. In result, we intend to publish a white-paper that systematically analyzes the existing traceability gap based on the outcome of the four focus groups. Furthermore, the workshop has initiated collaborations and potential research projects between previously separate areas with the potential of significant impact.

[43] *Mind the gap: Assessing the conformance of software traceability to relevant guidelines*, P. Rempel, P. Mäder, T. Kuschke, and J. Cleland-Huang, Proc. of the 36th Int'l Conf. on Software Engineering (ICSE'14), http://dx.doi.org/10.1145/2568225.2568290.

7.26 Theory and Practice of SAT Solving

Organizers: Armin Biere, Vijay Ganesh, Martin Grohe, Jakob Nordström, and Ryan Williams

Seminar No. 15171

Date: April 19–24, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.98

Participants: Erika Abraham, Albert Atserias, Gilles Audemard, Paul Beame, Christoph Berkholz, Olaf Beyersdorff, Armin Biere, Nikolaj S. Bjørner, Ilario Bonacina, Sam Buss, Amin Coja-Oghlan, Jessica Davies, Holger Dell, Jan Elffers, John Franco, Andreas Fröhlich, Vijay Ganesh, Alexandra Goultiaeva, Martin Grohe, Daniel Große, Edward A. Hirsch, Holger H. Hoos, Matti Järvisalo, Mikoláš Janota, Manuel Kauers, Oliver Kullmann, Massimo Lauria, Daniel Le Berre, Sharad Malik, João Marques-Silva, Nina Narodytska, Jakob Nordström, Yakau Novikau, Albert Oliveras, Pavel Pudlák, Alexander Razborov, Kristin Yvonne Rozier, Karem A. Sakallah, Martina Seidl, Laurent Simon, Niklas Sörensson, Stefan Szeider, Alasdair Urquhart, Allen Van Gelder, Marc Vinyals, Magnus Wahlström, Ryan Williams

This seminar brought together researchers working in the areas of applied SAT solving on the one hand, and in proof complexity and neighbouring areas of computational complexity theory on the other, in order to communicate new ideas, techniques, and analysis from both the practical and theoretical sides.

The goals of this endeavour are to better understand why modern SAT solvers work so efficiently for many large-scale real-world instances, and in the longer term to discover new strategies for SAT solving that could go beyond the present "conflict-driven clause-learning" paradigm and deliver substantial further gains in practical performance.

Topics of the Seminar

This seminar explored one of the most significant problems in all of mathematics and computer science, namely that of proving logic formulas. This is a problem of immense importance both theoretically and practically. On the one hand, it is believed to be intractable in general, and deciding whether this is so is one of the famous million dollar Clay Millennium Problems (the *P* vs. *NP* problem). On the other hand, today so-called SAT solvers are routinely and successfully used to solve large-scale real-world instances in a wide range of application areas (such as hardware and software verification, electronic design automation, artificial intelligence research, cryptography, bioinformatics, operations research, and railway signalling systems, just to name a few examples).

During the last 15–20 years, there have been dramatic – and surprising – developments in SAT solving technology that have improved real-world performance by many orders of magnitude. But perhaps even more surprisingly, the best SAT solvers today are still based on relatively simple methods from the early 1960s, searching for proofs in the so-called resolution proof system. While such solvers can often handle formulas with millions of variables, there are also known tiny formulas with just a few hundred variables that cause even the very best solvers to stumble. The fundamental question of when SAT solvers perform well or badly, and what underlying properties of the formulas influence SAT solver performance, remains very poorly understood. Other practical SAT solving issues, such as how to optimize memory management and how to exploit parallelization on modern multicore architectures, are even less well studied and understood from a theoretical point of view.

Another intriguing fact is that although other mathematical methods of reasoning are known that are much stronger than resolution in theory, in particular methods based on algebra and geometry, attempts to harness the power of such methods have failed to deliver any significant improvements in practical performance – indeed, such solvers often struggle even to match the performance of resolution-based solvers. And while resolution is a fairly well-understood proof system, even very basic questions about these stronger algebraic and geometric methods remain wide open.

We believe that computational complexity can shed light on the power and limitations on current and possible future SAT solving techniques, and that problems encountered in SAT solving can spawn interesting new areas in theoretical research. We see great potential for interdisciplinary research at the border between theory and practice in this area, and believe that more vigorous interaction between practitioners and theoreticians could have major long-term impact in both academia and industry.

Goals of the Seminar

A strong case can be made for the importance of increased exchange between the two fields of SAT solving on the one hand and proof complexity (and more broadly computational complexity) on the other. While the two areas have enjoyed some

exchanges, it seems fair to say that there has been relatively low level of interaction, given how many questions would seem to be of mutual interest. Below, we try to outline some such questions that served as motivation for organizing this seminar. We want to stress that this list is far from exhaustive, and in fact we believe one important outcome of the seminar was to stimulate the process of uncovering other questions of common interest.

What Makes Formulas Hard or Easy in Practice for Modern SAT Solvers?

The best SAT solvers known today are based on the DPLL procedure, augmented with optimizations such as conflict-driven clause learning (CDCL) and restart strategies. The propositional proof system underlying such algorithms, resolution, is arguably the most well-studied system in all of proof complexity.

Given the progress during the last decade on solving large-scale instances, it is natural to ask what lies behind the spectacular success of CDCL solvers at solving these instances. And given that there are still very small formulas that resist even the most powerful CDCL solvers, a complementary interesting question is if one can determine whether a particular formula is hard or tractable. Somewhat unexpectedly, very little turns out to be known about these questions.

In view of the fundamental nature of the SAT problem, and in view of the wide applicability of modern SAT solvers, this seems like a clear example of a question of great practical importance where the theoretical field of proof complexity could potentially provide useful insights. In particular, one can ask whether one could find theoretical complexity measures for formulas than would capture the practical hardness of these formulas in some nice and clean way. Besides greatly advancing our theoretical understanding, answering such a question could also have applied impact in the longer term by clarifying the limitations, and potential for further improvements, of modern SAT solvers.

Can Proof Complexity Shed Light on Crucial SAT Solving Issues?

Understanding the hardness of proving formulas in practice is not the only problem for which more applied researchers would welcome contributions from theoretical computer scientists. Examples of some other possible practical questions that would merit from a deeper theoretical understanding follow below.

- Firstly, we would like to study the question of memory management. One major concern for clause learning algorithms is to determine how many clauses to keep in memory. Also, once the algorithm runs out of the memory currently available, one needs to determine which clauses to throw away. These questions can have huge implications for performance, but are poorly understood.
- In addition to clause learning, the concept of restarts is known to have decisive impact on the performance on modern CDCL solvers. It would be nice to understand theoretically why this is so. The reason why clause learning increases efficiency greatly is clear – without it the solver will only generate so-called tree-like proofs, and tree-like resolution is known to be exponentially weaker than general resolution. However, there is still ample room for improvement of our understanding of the role of restarts and what are good restart strategies.
- Given that modern computers are multi-core architectures, a highly topical question is whether this (rather coarse-grained) parallelization can be used to speed up SAT solving. Our impression is that this is an area where much practical work is being carried out, but where comparatively little theoretical study has been done. Thus, the first step here would consist of understanding what are the right questions to ask and coming up with a good theoretical framework for investigating them. While there are some successful attempts in parallelizing SAT, obtained speed-ups are rather modest. This is a barrier for further adoption of SAT technology already today and will be become a more substantial problem as thousands of cores and cloud computing are becoming the dominant computing platforms. A theoretical understanding on how SAT can be parallelized will be essential to develop new parallelization strategies to adapt SAT to this new computing paradigm.

Can we build SAT Solvers based on Stronger Proof Systems than Resolution?

Although the performance of modern CDCL SAT solvers is impressive, it is nevertheless astonishing, not to say disappointing, that the state-of-the-art solvers are still based on simple resolution. Resolution lies very close to the bottom in the hierarchy of propositional proof systems, and there are many other proof systems based on different forms of mathematical reasoning that are known to be strictly stronger. Some of these appear to be natural candidates for serving as a basis for stronger SAT solvers than those using CDCL.

In particular, proof systems such as polynomial calculus (based on algebraic reasoning) and cutting planes (based on geometry) are known to be exponentially more powerful than resolution. While there has been some work on building SAT solvers on top of these proof systems, progress has been fairly limited. As part of the seminar, we invited experts on algebraic and geometric techniques to discuss what the barriers are that stops us from building stronger algebraic or geometric SAT solvers, and what is the potential for future improvements. An important part of this work would seem to be to gain a deeper theoretical understanding of the power and limitations of these proof methods. Here there are a number of fairly long-standing open theoretical questions. At the same time, only in the last couple of years proof complexity has made substantial progress, giving hope that the time is ripe for decisive break-throughs in these areas.

Organization of the Seminar

The scientific program of the seminar consisted of 26 talks. Among these there were five 80-minute tutorials on core topics of the seminar:

- proof complexity (Paul Beame),
- conflict-driven clause learning (CDCL) SAT solvers (João Marques-Silva),
- proof systems connected to SAT solving (Sam Buss),
- preprocessing and inprocessing (Matti Järvisalo),
- SAT and SMT (Nikolaj Bjørner).

Throughout, the tutorials were well-received as a means of introducing the topics and creating a common frame of reference for participants from the different communities.

There were also nine slightly shorter survey talks of 50 minutes which were intended to give overviews of a number of important topics for the seminar:

- semialgebraic proof systems (Albert Atserias),
- pseudo-Boolean constraints and CDCL (Daniel Le Berre),
- Gröbner bases (Manuel Kauers),
- SAT-enabled verification of state transition systems, (Karem Sakallah),
- SAT and computational complexity (Ryan Williams)

- the (strong) exponential time hypothesis and consequences (Ryan Williams),
- SAT and parameterized complexity (Stefan Szeider),
- QBF solving (Nina Narodytska),
- random satisfiability (Dimitris Achlioptas).

Most tutorials and survey talks were scheduled early in the week, to create a conducive atmosphere for collaboration on open problems later in the week. The rest of the talks were 25-minute presentations on recent research of the participants. The time between lunch and afternoon coffee was left for self-organized collaborations and discussions, and there was no schedule on Wednesday afternoon.

Based on polling of participants before the seminar week, it was decided to have an open problem session on Monday evening, and on Wednesday evening there was a panel discussion. The organizing committee also considered the option of having a poster session to give more researchers the opportunity to present recent research results, but the feedback in the participant poll was negative and so this idea was dropped.

Fig. 7.9
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.27 Challenges and Trends in Probabilistic Programming

Organizers: Gilles Barthe, Andrew D. Gordon, Joost-Pieter Katoen, and Annabelle McIver

Seminar No. 15181

Date: April 26–30, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.123

Participants: Christel Baier, Gilles Barthe, Johannes Borgström, Michael Carbin, Aleksandar Chakarov, Ugo Dal Lago, Alessandra Di Pierro, Jason Eisner, Yuan Feng, Luis Maria Ferrer Fioriti, Cédric Fournet, Cameron Freer, Marco Gaboardi, Andrew D. Gordon, Friedrich Gretz, Johannes Hölzl, Chung-Kil Hur, Benjamin Kaminski, Joost-Pieter Katoen, Stefan Kiefer, Angelika Kimmig, Boris Köpf, Pasquale Malacaria, Vikash Mansinghka, Piotr Mardziel, Annabelle McIver, Joel Ouaknine, Catuscia Palamidessi, David Parker, Avi Pfeffer, Tahiry Rabehaja, Sriram K. Rajamani, Norman Ramsey, Chung-chieh Shan, Alexandra Silva, Sameer Singh, Geoffrey Smith, Andreas Stuhlmüller, Frank Wood, Mingsheng Ying, Lijun Zhang

Probabilistic programming languages

Probabilistic programs are programs, written in languages like C, Java, LISP, or ML, with two added constructs: (1) the ability to draw values at random from probability distributions, and (2) the ability to condition values of variables in a program through observations. A variety of probabilistic programming languages have been defined such as `Church`, `Infer.NET`, and `IBAL`. `Church` is based on the Lisp model of the lambda calculus, containing pure Lisp as its deterministic subset, whereas `Infer.NET` is a Microsoft developed language akin to `C#` and compiles probabilistic programs into inference code[44]. Probabilistic programs can be used for modelling complex phenomena from biology and social sciences. By doing so, we get the benefits of programming languages (rigorous semantics, execution, testing and verification) to these problem domains. More than a decade ago, McIver and Morgan defined a probabilistic programming language in the style of Dijkstra's guarded command language, referred to as `pGCL`. Besides the usual language constructs in Dijkstra's `GCL` such as non-deterministic choice, it features a probabilistic choice where the probability distribution may be parametric. For instance, the assignment $\mathtt{x} \mathrel{+}= 1\ [p]\ \mathtt{skip}$ increments the variable `x` by one with probability p, and keeps the value of `x` unchanged with probability $1-p$, where p is an unknown real value from the range $[0, 1]$. Quantum programming languages such as `qGCL` and a quantum extension of `C++` are also related, as their operational semantics is typically a probabilistic model so as to model the effect of measurements on the quantum state.

The importance of probabilistic programming The applications of probabilistic programs mainly lie in four domains: (1) machine learning, (2) security, (3) randomised algorithms, and – though to a somewhat lesser extent – (4) quantum computing. Whereas the application in the field of randomised algorithms is evident, let us briefly describe the importance for the other three fields.

Machine learning. A Bayesian generative model consists of a prior distribution over some parameters, together with a sampling distribution (or likelihood) that predicts outputs of the model given its inputs and parameters. Bayesian inference in machine learning consists of training such a model to infer the posterior distribution of the parameters and hence to make predictions. In the probabilistic programming approach to Bayesian inference, the user simply writes the prior and sampling distributions as probabilistic programs, and relies on a compiler to generate code to perform inference and make predictions. Such compilers often operate by considering the program as defining a probabilistic graphical model. Graphical models were pioneered by Judea Pearl and others, and are extensively described in the comprehensive text by Koller and Friedman (2009). They are widely used in statistics and machine learning, with diverse application areas including speech recognition, computer vision, biology, and reliability analysis. Probabilistic graphical models allow specification of dependences between random variables via generative models, as well as conditioning of random variables using phenomena or data observed in the real world. A variety of inference algorithms have been developed to analyse and query such models, e.g., Gibbs sampling methods, Metropolis-Hastings and belief propagation. The probabilistic programming approach has seen growing interest within machine learning over the last 10 years and it is believed – see http://probabilistic-programming.org/wiki/Home – that this approach within AI has the potential

[44] For academic use, it is free to use: http://research.microsoft.com/infernet.

to fundamentally change the way that community understands, designs, builds, tests and deploys probabilistic systems.

Security. Ever since Goldwasser and Micali – recipients of the ACM Turing Award in 2013 – introduced probabilistic encryption, probability has played a central role in cryptography: virtually all cryptographic algorithms are randomized, and have probabilistic security guarantees. Similarly, perturbing outputs with probabilistic noise is a standard tool for achieving privacy in computations; for instance, differential privacy achieves privacy-preserving data-mining using probabilistic noise. Cryptographic algorithms and differentially private algorithms are implemented as probabilistic programs; more singularly, one common approach for reasoning about these algorithms is using the code-based game-based approach, proposed by Bellare and Rogaway, in which not only the algorithms, but also their security properties and the hardness properties upon which their security relies, are expressed as probabilistic programs , and can be verified using (a relational variant of) Hoare logic. This code-based approach is key to recent developments in verified cryptography. Quantitative information flow is another important field in security where probabilistic programs and models play an important role. Here, the key question is to obtain quantitative statements about the leakage of certain information from a given program.

Quantum computing. Quantum programs are used to describe quantum algorithms and typically are quantum extensions of classical while-programs. Whereas in classical computation, we use a type to denote the domain of a variable, in quantum computation, a type is the state space of a quantum system denoted by some quantum variable. The state space of a quantum variable is the Hilbert space denoted by its type. According to a basic postulate of quantum mechanics, the unique way to acquire information about a quantum system is to measure it. Therefore, the essential ingredient in a quantum program is the ability to perform measurements of quantum registers, i.e., finite sequences of distinct quantum variables. The state space of a quantum register is the tensor product of the state spaces of its quantum variables. In executing the statement `measure M[q]; S`, quantum measurement M will first be performed on quantum register q, and then a sub-program S will be selected to be executed next according to the outcome of the measurement. The essential difference between a measurement statement and a classical conditional statement is that the state of program variables is changed after performing the measurement. As the outcome of a measurement is probabilistic, quantum programs are thus inherently probabilistic.

Program analysis. On the other hand, there is a recent rapidly growing trend in research on probabilistic programs which is more in line with traditional programming languages. This focuses on aspects such as efficient compilation, static analysis, program transformations, and program verification. To mention a few, Cousot *et al.* recently extended the framework of abstract interpretation to probabilistic programs (2012), Gordon *et al.* introduced `Tabular`, a new probabilistic programming language (2014), Di Pierro *et al.* apply probabilistic static analysis (2010), Rajamani, Gordon *et al.* have used symbolic execution to perform Bayesian reasoning on probabilistic programs with loops (2013), Katoen, McIver *et al.* have developed invariant synthesis technique for linear probabilistic programs (2010), and Geldenhuys *et al.* considered probabilistic symbolic execution (2012).

Achievements of this seminar

The objective of the seminar was a to bring together researchers from four separate (but related) communities to learn from each other, with the expectation that a better understanding between these communities would open up new opportunities for research and collaboration.

Participants attending the seminar represented all four themes of the original proposal: machine learning, quantitative security, (probabilistic) program analysis and quantum computing. The programme consisted of both tutorials and presentations on any topic within these themes. The tutorials provided a common ground for discussion, and the presentations gave insight into the current state of an area, and summarised the challenges that still remain. The tutorial topics were determined by consulting the participants prior to the seminar by means of a questionnaire.

Although the programme was primarily constructed around the tutorials and standard-length presentations (each around 30 minutes), the organisers made sure that time was always available for short, impromptu talks (sometimes of only 5 minutes) where participants were able to outline a relevant challenge problem or to draw attention to a new research direction or connection that had become apparent during the meeting.

This open forum for exploring links between the communities has led to the following specific achievements:

1. An increased understanding between the disciplines, especially between program verification and probabilistic programming.
2. A demonstration that the mathematical models for reasoning about machine learning algorithms and quantitative security are very similar, but that their objectives are very different. This close relationship at a foundational level suggests theoretical methods to tackle the important challenge of understanding privacy in a data mining context.
3. Evidence that probabilistic programming, analysis and verification of probabilistic programs, can have a broad impact in the design of emerging infrastructures, such as software-defined networks.

The feedback by the participants was very positive, and it was encouraged to organise a workshop or similar event in the future to foster the communication between the different communities, in particular between program verification and probabilistic programming.

We were aware of many new conversations between researchers inspired by the formal talks as well as the mealtime discussions. Already at least one paper (see below) with content inspired by the meeting is accepted for publication, and we are aware of several other new lines of work.

Acknowledgement. The organisers thank Benjamin Kaminski for his support in several organisational issues.

References

1 A. Ścibior, Z. Ghahramani and A. D. Gordon. *Practical Probabilistic Programming with Monads.* ACM SIGPLAN Haskell Symposium 2015, Vancouver, Canada, 3–4 September 2015.

7.28 Qualification of Formal Methods Tools

Organizers: Darren Cofer, Gerwin Klein, Konrad Slind, and Virginie Wiels

Seminar No. 15182

Date: April 26–29, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.4.142

Participants: June Andronick, Rob Arthan, Jasmin Christian Blanchette, Sandrine Blazy, Matteo Bordin, Darren Cofer, David Cok, Rémi Delmas, Michael Dierkes, Eric Engstrom, Gerwin Klein, Ramana Kumar, Mark Lawford, Xavier Leroy, Stefan Leue, Alain Mebsout, Stephan Merz, Cesar A. Munoz, Magnus Myreen, Scott Owens, Marc Pantel, Markus Pister, Werner Schütz, Konrad Slind, Nick Tudor, Lucas Wagner, Michael W. Whalen, Virginie Wiels

Motivation and objectives

Dagstuhl Seminar 13051, *Software Certification: Methods and Tools*, convened experts from a variety of software-intensive domains (automotive, aircraft, medical, nuclear, and rail) to discuss software certification challenges, best practices, and the latest advances in certification technologies. One of the key challenges identified in that seminar was tool qualification. Tool qualification is the process by which certification credit may be claimed for the use of a software tool. The purpose of tool qualification is to provide sufficient confidence in the tool functionality so that its output may be trusted. Tool qualification is, therefore, a significant aspect of any certification effort. Seminar participants identified a number of needs in the area of formal methods tool qualification. Dagstuhl Seminar 15182 *Qualification of Formal Methods Tools*, was organized to address these needs.

Software tools are used in development processes to automate life cycle activities that are complex and error-prone if performed by humans. The use of such tools should, in principle, be encouraged from a certification perspective to provide confidence in the correctness of the software product. Therefore, we should avoid unnecessary barriers to tool qualification which may inadvertently reduce the use of tools that would otherwise enhance software quality and confidence.

Most software tools are not used in isolation, but are used as part of a complex tool chain requiring significant integration effort. In general, these tools have been produced by different organizations. We need to develop better and more reliable methods for integrating tools from different vendors (including university tools, open source tools, and commercial tools).

A given software tool may be used in different application domains having very different requirements for both certification and tool qualification. Furthermore, the methods and standards for tool development varies across domains. Consistent qualification requirements across different domains would simplify the process.

Despite the additional guidance provided for the avionics domain in recently published standards (DO-178C, DO-330, and DO-333), there are still many questions to be addressed. For one thing, most practicing engineers are unaware of how to apply different categories of formal verification tools. Even within a particular category, there are a wide variety of tools, often based on fundamentally different approaches, each with its own strengths and weaknesses.

If formal verification is used to satisfy DO-178C objectives, DO-333 requires the applicant to provide evidence that the underlying method is sound, i.e., that it will never assert something is true when it is actually false, allowing application software errors to be missed that should have been detected. Providing an argument for the soundness of a formal verification method is highly dependent on the underlying algorithm on which the method is based. A method may be perfectly sound when used one way on a particular type of problem and inherently unsound when used in a different way or on a different type of problem. While these issues may be well understood in the research community, they are not typically collected in one place where a practitioner can easily find them. It is also not realistic to expect avionics developers to be able to construct an argument for the soundness of a formal method without help from experts in the field.

At the same time, it is also important to not make the cost of qualification of formal methods tools so great as to discourage their use. While it is tempting to hold formal verification tools to a higher standard than other software tools, making their qualification unnecessarily expensive could do more harm than good.

The objectives of this Dagstuhl Seminar were to

- investigate the sorts of assurances that are necessary and

appropriate to justify the application of formal methods tools throughout all phases of design in real safety-critical settings,
- discuss practical examples of how to qualify different types of formal verification tools, and
- explore promising new approaches for the qualification of formal methods tools for the avionics domain, as well as in other domains.

Accomplishments

Qualification is not a widely understood concept outside of those industries requiring certification for high-assurance, and different terminology is used in different domains. The seminar was first a way of sharing knowledge from certification experts so that formal methods researchers could better understand the challenges and barriers to the use of formal methods tools.

The seminar also included presentations from researchers who have developed initial approaches to address qualification requirements for different classes of formal methods tools. We were especially interested in sharing case studies that are beginning to address tool qualification challenges. These case studies include tools based on different formal methods (model checking, theorem proving, abstract interpretation).

As a practical matter, we focussed much of our discussion on the aerospace domain since there are published standards addressing both formal methods and tool qualification for avionics software. The seminar also included researchers from other domains (nuclear, railway) so we could better understand the challenges and tool qualification approaches that are being discussed in those domains.

We managed to bridge a lot of the language between the certification domains, mostly railway, avionics, and nuclear, and bits of automotive, and related the qualification requirements to each other. Some of the otherwise maybe less stringent schemes (e.g. automotive) can end up having stronger qualification requirements, because formal methods are not specifically addressed in them. There is some hope that DO-333 might influence those domains, or be picked up by them in the future, to increase the use of FM tools which would increase the quality of systems.

For the academic tool provider side, we worked out and got the message across that tool qualification can be a lot easier and simpler than what we might strive for academically, and discussed specific tools in some detail, clarifying what would be necessary for a concrete qualification. Finally, we also investigated tool architectures that make tools easier to qualify (verification vs code generation).

7.29 Compositional Verification Methods for Next-Generation Concurrency

Organizers: Lars Birkedal, Derek Dreyer, Philippa Gardner, and Zhong Shao

Seminar No. 15191

Date: May 3–8, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.5.1

Participants: Lennart Beringer, Lars Birkedal, Andrea Cerone, Adam Chlipala, Karl Crary, Pedro Da Rocha Pinto, Thomas Dinsdale-Young, Mike Dodds, Alastair F. Donaldson, Cezara Dragoi, Derek Dreyer, Xinyu Feng, Philippa Gardner, Alexey Gotsman, Aquinas Hobor, Jan Hoffmann, Chung-Kil Hur, Bart Jacobs, Cliff B. Jones, Ralf Jung, Eric Koskinen, Neel Krishnaswami, Ori Lahav, Hongjin Liang, Paul McKenney, Maged M. Michael, Peter Müller, Aleksandar Nanevski, Kyndylan Nienhuis, Scott Owens, Frank Pfenning, Jean Pichon-Pharabod, Francois Pottier, Shaz Qadeer, Azalea Raad, John Reppy, Noam Rinetzky, Claudio Russo, Ilya Sergey, Peter Sewell, Zhong Shao, Kasper Svendsen, Joseph Tassarotti, Viktor Vafeiadis, Martin T. Vechev, Nobuko Yoshida

One of the major open problems confronting software developers today is how to cope with the complexity of reasoning about large-scale concurrent programs. Such programs are increasingly important as a means of taking advantage of parallelism in modern architectures. However, they also frequently depend on subtle invariants governing the use of shared mutable data structures, which must take into account the potential interference between different threads accessing the state simultaneously. Just figuring out how to express such invariants at all has proven to be a very challenging problem; even more challenging is how to support *local* reasoning about such invariants, i.e., confining the reasoning about them to only the components of the program that absolutely need to know about them.

Fortunately, we are now at a point where verification research has produced the critical foundations needed to tackle this problem: namely, *compositional* methods, which exploit the inherently modular structure of realistic concurrent programs in order to decompose verification effort along module boundaries. Fascinatingly, a variety of different but related compositional methods have been developed contemporaneously in the last several years:

- **Separation logics:** Separation logic was developed initially as a generalization of Hoare logic – supporting local, compositional reasoning about sequential, heap-manipulating programs – and much of the early work on separation logic has been successfully incorporated into automated verification tools like Smallfoot [2], SLAyer [3], Abductor [6], etc., scaling to handle millions of lines of code. Recently, there have been a series of breakthroughs in adapting separation logic to handle concurrent programs as well. Concurrent separation logic [17] provides course-grained local reasoning about concurrent programs; combining this local reasoning with rely-guarantee reasoning [26] provides fine-grained concurrent reasoning; intertwining abstraction with local reasoning enables a client to reason about the use of a set module [8] without having to think about the underlying implementation using lists or concurrent B-trees; and, very recently, all this has been extended to account for higher-order programs as well [21].
- **Kripke models:** There is a long line of work on the use of semantic models like *Kripke logical relations* [1, 9] (and more recently *bisimulations* [19, 20]) for proving observational equivalence of programs that manipulate local state. Observational equivalence is useful not only for establishing correctness of program transformations (e.g., in compiler certification) but also as a verification method in its own right (e.g., one can prove that a complex but efficient implementation of an ADT is equivalent to a simple but inefficient reference implementation). However, it is only in the last few years that such models have been generalized to account for the full panoply of features available in modern languages: higher-order state, recursion, abstract types, control operators, and most recently concurrency, resulting in some of the first formal proofs of correctness of sophisticated fine-grained concurrent algorithms in a higher-order setting [1, 9, 23]. These advances have come about thanks to the development of more elaborate Kripke structures for representing invariants on local state.
- **Hoare type theory:** Dependent type theory provides a very expressive compositional verification system for higher-order functional programs, so expressive that types can characterize full functional correctness. Traditionally, however, dependent type theories were limited to verification of *pure* programs. Recent work on Hoare type theory (HTT) [15] has shown how to integrate effects into dependent type theory by incorporating Hoare triples as a new primitive type, and prototypes of HTT have been implemented in Coq [7, 16], allowing for imperative programs to be verified mechanically as they are

being written. Moreover, first steps of extending HTT with concurrency have recently been taken [14], thus giving hope for a potential future integration of design and verification for higher-order concurrent programs.

All in all, the field of modular concurrency verification is highly active, with groundbreaking new developments in these and other approaches coming out every year. Particularly fascinating is the appearance of deep connections between the different methods. There are striking similarities, for instance, between the advanced Kripke structures used in recent relational models of higher-order state and the semantic models underlying recent concurrent separation logics.

Nevertheless, there are a number of ways in which the advanced models and logics developed thus far are still, to be honest, in their infancy. Most of these approaches, for example, have only been applied to the verification of small, self-contained ADTs and have not yet been scaled up to verify large-scale modular concurrent programs. Moreover, even the most state-of-the-art compositional methods do not yet account for a number of the essential complexities of concurrent programming as it is practiced today, including:

- **Weak memory models:** The vast majority of state-of-the-art compositional verification methods are proved sound with respect to an operational semantics that assumes a sequentially consistent memory model. However, modern hardware implements weak memory models that allow for many more reorderings of basic operations. Thus there is a clear gap between the verification theory and practice that needs to be filled (for efficiency reasons we, of course, do not want to force programmers/compilers to insert enough memory fence operations to make the hardware behave sequentially consistent). This problem has been known for the last decade, but it is only in the last year or two that formal descriptions of the behavior of programming languages with weak memory models have been developed. Given this foundation, we should now be able to make progress on extending compositional verification methods to weak memory models.
- **Higher-order concurrency:** Higher-order functional abstraction is an indispensable feature of most modern, high-level programming languages. It is also central to a variety of concurrent programming idioms, both established and nascent: work stealing [4], Concurrent ML-style events [18], concurrent iterators [13], parallel evaluation strategies [22], STM [11], reagents [24], and more. Yet, only a few existing logics have been proposed that even attempt to account for higher-order concurrency [12,14,21], and these logics are just first steps – for example, they do not presently account for sophisticated "fine-grained" concurrent ADTs. Verification of higher-order concurrent programs remains a largely open problem.
- **Generalizing linearizability:** Sophisticated concurrent data structures often use fine-grained synchronization to maximize the possibilities for parallel access. The classical correctness criterion for such fine-grained data structures is *linearizability*, which ensures that every operation has a linearization point at which it appears (to clients) to atomically take effect. However, existing logics do not provide a way to exploit linearizability directly in client-side reasoning, and moreover the notion does not scale naturally to account for operations (such as higher-order iterators) whose behavior is not semantically atomic. Recently, researchers have started to investigate alternative approaches, based on *contextual refinement* [10, 23]. And methods for reasoning about operations with multiple linearizability points are also being developed.
- **Liveness properties:** Synchronization of concurrent data structures can also affect the progress of the execution of the client threads. Various progress properties have been proposed for concurrent objects. The most important ones are wait-freedom, lock-freedom and obstruction-freedom for non-blocking implementations, and starvation-freedom and deadlock-freedom for lock-based implementations. These properties describe conditions under which method calls are guaranteed to successfully complete in an execution. Traditional definitions (which are quite informal) of these progress properties are difficult to use in modular program verification because they fail to describe how the progress properties affect clients. It is also unclear how existing separation logics, which were primarily designed for proving partial correctness, can be adapted to prove progress properties. Recently, researchers have started to combine *quantitative reasoning* of resource bounds with separation logics, which offer new possibilities for verifying both safety and liveness properties in a single framework.

Grappling with these kinds of limitations is essential if our verification technology is to be relevant to real-world programs running on modern architectures, and as such it poses exciting new research questions that we as a community are just beginning to explore.

In this seminar, we brought together a wide variety of researchers on concurrency verification, as well as leading experts on concurrent software development in both high- and low-level languages. The goal was to facilitate a stimulating interchange between the theory and practice of concurrent programming, and thereby foster the development of compositional verification methods that can scale to handle the realities of next-generation concurrency.

Among the concrete research challenges investigated in depth during the seminar are the following:

- What are good ways of reasoning about weak memory models? It should be possible to reason about low-level programs that exploit weak memory models (e.g., locks used inside operating systems) but also to reason at higher levels of abstractions for programs that use sufficient locking.
- What is the best way to define a language-level memory model that is nevertheless efficiently implementable on modern hardware. C11 is the state of the art, but it is flawed in various ways, and we heard about a number of different ways of possibly fixing it.
- What is the best way to mechanize full formal verification of concurrent programs, using interactive proof assistants, such as Coq.
- How can we adapt existing and develop new compositional techniques for reasoning about liveness properties of concurrent programs? Can we apply quantitative techniques to reduce the proof of a liveness property to the proof of a stronger safety property? Also, recent work on rely-guarantee-based simulation can prove linearizability of a sophisticated concurrent object by showing the concurrent implementation is a contextual refinement of its sequential specification. We would hope that similar techniques can be used to prove progress properties as well.
- Only recently have researchers begun to propose logics and models for higher-order concurrency [21, 23]. What are the right concurrency abstractions for higher-order concurrent programming idioms as diverse as transactional memory [11], Concurrent ML [18], joins [25], and reagents [24], among

others? What is the best way to even specify, let alone verify, programs written in these idioms, and are there unifying principles that would apply to multiple different idioms?

- Most verification work so far has focused on shared-memory concurrency, with little attention paid to message-passing concurrency (except for some recent work on verifying the C♯ joins library). Can the models and logics developed for the former be carried over usefully to the latter, and what is the connection (if any) with recent work on proof-theoretic accounts of session types [5]? Can session types help to simplify reasoning about some classes of concurrent programs, e.g., those that only involve some forms of message passing and not full shared memory?
- A number of recent Kripke models and separation logics have employed *protocols* of various forms to describe the invariants about how the semantic state of a concurrent ADT can evolve over time. But different approaches model protocols differently, e.g., using enriched forms of state transition systems vs. partial commutative monoids. Is there a canonical way of representing these protocols formally and thus better understanding the relationship between different proof methods?
- There seem to be tradeoffs between approaches to concurrency verification based on Hoare logic vs. refinement (unary vs. relational reasoning), with the former admitting a wider variety of formal specifications but the latter offering better support for reasoning about atomicity. Consequently, a number of researchers are actively working on trying to combine both styles of reasoning in a unified framework. What is the best way to do this?
- To what extent do we need linearizability to facilitate client-side reasoning? Is it possible in many cases for clients to rely on a much weaker specification ? And which ways are there to formalize looser notions, e.g. where there are multiple linearization points?
- Now that we are finally developing logics and models capable of verifying realistic concurrent algorithms, can we abstract away useful proof patterns and automate them? What is needed in order to integrate support for concurrent invariants into automated verification tools like SLAyer and Abductor?

These different challenges were discussed through talks and discussions by participants, see the list of talk abstracts below.

References

1 A. Ahmed, D. Dreyer, and A. Rossberg. State-dependent representation independence. In *POPL*, 2009.

2 J. Berdine, C. Calcagno, and P. W. O'Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *FMCO*, 2005.

3 J. Berdine, B. Cook, and S. Ishtiaq. Slayer: Memory safety for systems-level code. In *CAV*, 2011.

4 R. D. Blumofe, C. F. Joerg, B. C. Kuszmaul, C/ E. Leiserson, K. H. Randall, and Y. Zhou. Cilk: An efficient multithreaded runtime system. *JPDC*, 37(1):55–69, August 1996.

5 L. Caires and F. Pfenning. Session types as intuitionistic linear propositions. In *CONCUR*, 2010.

6 C. Calcagno, D. Distefano, P. W. O'Hearn, and H. Yang. Compositional shape analysis by means of bi-abduction. *Journal of the ACM*, 58(6), 2011.

7 A. Chlipala, G. Malecha, G. Morrisett, A. Shinnar, and R. Wisnesky. Effective interactive proofs for higher-order imperative programs. In *ICFP*, 2009.

8 T. Dinsdale-Young, M. Dodds, P. Gardner, M. Parkinson, and V. Vafeiadis. Concurrent abstract predicates. In *ECOOP*, June 2010.

9 D. Dreyer, G. Neis, and L. Birkedal. The impact of higher-order state and control effects on local relational reasoning. *Journal of Functional Programming*, 22(4&5):477–528, August 2012.

10 I. Filipović, P. W. O'Hearn, N. Rinetzky, and H. Yang. Abstraction for concurrent objects. In *ESOP*, 2009.

11 T. Harris, S. Marlow, S. Peyton-Jones, and M. Herlihy. Composable memory transactions. In *PPoPP*, 2005.

12 A. Hobor, A. W. Appel, and F. Zappa Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, 2008.

13 D. Lea. The java.util.concurrent ConcurrentHashMap.

14 R. Ley-Wild and A. Nanevski. Subjective auxiliary state for coarse-grained concurrency. In *POPL*, 2013.

15 A. Nanevski, G. Morrisett, and L. Birkedal. Hoare type theory, polymorphism and separation. *JFP*, 18(5-6):865–911, 2008.

16 A. Nanevski, V. Vafeiadis, and J. Berdine. Structuring the verification of heap-manipulating programs. In *POPL*, 2010.

17 P. W. O'Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science*, 375(1–3):271–307, May 2007.

18 J. H. Reppy. *Higher-order concurrency*. PhD thesis, Cornell University, 1992.

19 D. Sangiorgi, N. Kobayashi, and E. Sumii. Environmental bisimulations for higher-order languages. In *LICS*, 2007.

20 E. Sumii. A complete characterization of observational equivalence in polymorphic λ-calculus with general references. In *CSL*, 2009.

21 K. Svendsen, L. Birkedal, and M. Parkinson. Modular reasoning about separation of concurrent data structures. In *ESOP*, 2013.

22 P. W. Trinder, K. Hammond, H.-W. Loidl, and S. L. Peyton Jones. Algorithm + strategy = parallelism. *JFP*, 8(1):23–60, January 1998.

23 A. Turon, J. Thamsborg, A. Ahmed, L. Birkedal, and D. Dreyer. Logical relations for fine-grained concurrency. In *POPL*, 2013.

24 A. Turon. Reagents: expressing and composing fine-grained concurrency. In *PLDI*, 2012.

25 A. Turon and C. Russo. Scalable join patterns. In *OOPSLA*, 2011.

26 V. Vafeiadis. *Modular fine-grained concurrency verification*. PhD thesis, University of Cambridge, 2008.

Fig. 7.10
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.30 The Message in the Shadow: Noise or Knowledge?

Organizers: Roberto Casati, Patrick Cavanagh, and Paulo E. Santos

Seminar No. 15192

Date: May 3–8, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.5.24

Participants: Juliano Beraldo, Mehul Bhatt, Francesca Bizzarri, Roberto Casati, Patrick Cavanagh, Marcos Cesar Danhoni Neves, Hannah Dee, Valquiria Fenelon Pereina, Frederic Fol Leymarie, Christian Freksa, John Kennedy, John O'Dea, Ann Marie Raynal, Paulo E. Santos, Carl Schultz, William Sharpe, Jakob Suchan, Koichi Toyama, Barbara Tversky, Maarten Wijntjes

The seminar "The Message in the Shadow: Noise or Knowledge?" brought together researchers from the various disciplines involved in investigating the problem of understanding the perception of shadows (both in biological and in artificial systems) as well as art historians and artists involved in the study or in the manipulation of shadows in art pieces. The nationalities of the seminar participants were as varied as the disciplines involved its central theme; from the 20 attendees there were 4 that came from Brazil, 4 from Germany, 1 from the Emirates, 2 from France, 2 from the UK, 1 from Canada, 3 from the US, 1 from the Netherlands and 2 from Japan.

The small size of the seminar helped to create a friendly atmosphere, in which every participant had time and space to engage in discussions with every other, and every one had an equal amount of time to present his/her ideas, independently of the career stage the participant was in.

The dynamics of the seminar was as follows: every participant that had an interest in presenting a talk was allocated a 20 minute slot, followed by a 10 min discussion period, during the mornings (from 9 to 11am). The talks where distributed into 4 tracks, one for each day of the week: Psychology (Monday), Artificial Intelligence and Computer Vision (Tuesday), Art and Rendering (Wednesday), Architecture and Spatial Reasoning (Thursday). The titles of the talks given, per track, are cited as follows (the related abstracts are listed in the next section):

Psychology

- Patrick Cavanagh. What does vision know about shadows?
- John Kennedy, Shape-from-shadow polarity
- John O'Dea, Do shadows make surfaces look dark?
- Marteen Wijntjes, Perception of shadows in paintings

Artificial Intelligence and Computer Vision

- Hannah M. Dee, Why does computer vision find shadows so problematic?
- Paulo E. Santos, Shadows in AI and Robotics
- Frederick Fol Leymarie, On medialness-based shape representation: recent developments and food for thought
- Ann Marie Raynal, Leveraging the Information in the Shadows of Synthetic Aperture Radar

Art and Rendering

- Koichi Toyama, The systematic introduction of Chiaroscuro in 15th century Florence and the symbolic shadow in Sienese Painting
- William Sharpe, Shadow Messages in the arts
- Marcos Danhoni, Shadows on the moon and the sun by Cigoliand Galileo: The Copernican planetarium inside the Paolina´s Chapel of Santa Maria Maggiore
- Roberto Casati, X-From-Shadow: There is still room at the bottom
- Koichi Toyama, Un-naturalistic painting and the lack of shadow: History of shadow in 18th- 19th century Japanese paintings and woodblock prints

Architecture and Spatial Reasoning

- Barbara Tversky, Can uses of shadows in language and art inform perception of shadows?
- Juliano Beraldo, Daylight metrics for building design
- Christian Freksa, Shadow and friends illuminate space
- Mehul Bhatt, Carl Schultz and Jakob Suchan, Grasping Objectified Shadows

Working Groups

At the end of the morning sessions, discussions were conducted in which the ideas presented during the talks served as inspiration for the conception of research statements. Some of these statements were selected to be discussed during the break out session that occurred during the Monday and Tuesday afternoons. The main questions discussed are presented below:

- Information about the light-source contained in shadows: there is a number of features from the light source that is present in the shadow of an object (for instance: the number of sources, the localisation, the shape) but much of this information is not used by the perceptual system. The question of the evolutionary advantages of this selective use of the information content of shadows was discussed and also the possibilities for a computer system to explore it fully;
- Mooney Faces and Shadows: To test people's vision, Craig Mooney devised two-tone pictures of faces. In Mooney faces, some parts are strongly illuminated, others are in deep shadow. His pictures were static. Motion helps vision find the faces. Mooney faces in negatives are hard to make out. Proper facial expression is lost. In outline, they are equally uninterpretable. Adding a dark line to the border of a positive Mooney face can drop recognition to the level of a negative. Motion helps, but still leaves the face looking cartoonish and flat. Often the line is taken as part of a profile. A light line border of a negative also leaves it cartoonish.
- Cross-disciplinary terminology for shadows: there is currently a non-consensual use of terms to refer to shadow issues (for instance, a caster is sometimes referred to as 'obtruder' or 'occluder'). This group proposed a tentative terminology that was later discussed with the other participants.
- Throwing away information. Shadows are used by the visual system to retrieve various spatial features of the scene, then discarded. The group discussed cognitive/computational mechanisms that may throw away shadows.
- Mereotopological formalisation of Eclipses. The group created a formalized version of the terminology used in describing the different phases of an Eclipse of the Sun. An amendment of the existing taxonomy was proposed.

At the end of the Monday session, artist Francesca Bizzarri showed some aspects of the art of shadow performance.

On Thursday afternoon the participants were directed to discuss possible collaborations, project proposals, and to devise conclusions (even if partial) to the various questions discussed during the previous days. Some of the results obtained in this session are listed below:

- Collaboration between S. Paulo and Bremen
- Online, real-time, Mooney face generator – A computer generated video by Dee, Kennedy and Casati, on the impairment of depth perception through the display of lines on moving Mooney faces, has been created and is visible at: https://www.youtube.com/watch?v=IuDNUz9RSuw
- Collaboration between Tokyo and New York (on art history)
- The foundation of a work group on terminology
- The projected publication of the mereotopological formalisation of Eclipses (Paris-Bremen-S. Paulo)
- Video displaying the phenomenon of the polarization of shadow (Casati and Cavanagh)

Finally, we discussed the future submission of a proposal for a special issue of the Journal Spatial Cognition and Computation (http://www.tandfonline.com/toc/hscc20/current) with the themes of the seminar and the organisation of a follow-up event in 2017 related to these ideas. Our proposal for a special issue of the Journal Spatial Cognition and Computation was accepted by the journal editors in June, 2015 (the *call for papers* will be advertised in the second half of this year).

7.31 Cross-Lingual Cross-Media Content Linking: Annotations and Joint Representations

Organizers: Alexander G. Hauptmann, James Hodson, Juanzi Li, Nicu Sebe, and Achim Rettinger

Seminar No. 15201

Date: May 10–13, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.5.43

Participants: Xavier Carreras, Dubravko Culibrk, John Davies, Tiansi Dong, Anastassia Fedyk, Blaz Fortuna, Marko Grobelnik, Alexander G. Hauptmann, James Hodson, Estevam R. Hruschka, Bea Knecht, Juanzi Li, Dunja Mladenic, Aditya Mogadala, Chong-Wah Ngo, Blaz Novak, Stefano Pacifico, Achim Rettinger, Evan Sandhaus, Nicu Sebe, Alan Smeaton, Rudi Studer, Jie Tang, Andreas Thalhammer, Eduardo Torres Schumann, Volker Tresp, Christina Unger, Michael Witbrock, Lexing Xie, Lei Zhang

Different types of content belonging to multiple modalities (text, audio, video) and languages are generated from various sources. These sources either broadcast information on channels like TV and News or allow collaboration in social media forums. Often multiple sources are consumed in parallel. For example, users watching TV tweeting their opinions about a show. This kind of consumption throw new challenges and require innovation in the approaches to enhance content search and recommendations.

Currently, most of search and content based recommendations are limited to monolingual text. To find semantic similar content across different languages and modalities, considerable research contributions are required from various computer science communities working on natural language processing, computer vision and knowledge representation. Despite success in individual research areas, cross-lingual or cross-media content retrieval has remained an unsolved research issue.

To tackle this research challenge, a common platform is provided in this seminar for researchers working on different disciplines to collaborate and identify approaches to find similar content across languages and modalities. After the group discussions between seminar participants, two possible solutions are taken into consideration:

1. *Building a joint space from heterogeneous data generated from different modalities to generate missing or to retrieve modalities. This is achieved through aligned media collections (like parallel text corpora). Now to find cross-media cross-lingual relatedness of the content mapped to a joint latent space, similarity measures can be used.*
2. *Another way is to build a shared conceptual space using knowledge bases (KB) like DBpedia etc for semantic annotation of concepts or events shared across modalities and languages. Entities are expressed in any channel, media type or language cam be mapped to a concept space in KB. Identifying a commonality between annotations can be used to find cross-media cross-lingual relatedness.*

Thus, implementing these solutions require a joint effort across research disciplines to relate the representations and to use them for linking languages and modalities. This seminar also aimed to build datasets that can be used as standard test bed and benchmark for cross-lingual cross-media content linking. Also, seminar was very well received by all participants. There was a common agreement that the areas of text, vision and knowledge graph should work more closely together and that each discipline would benefit from the other. The participants agreed to continue to work on two cross-modal challenges and discuss progress and future steps in a follow-up meeting in September at Berlin .

7

7.32 Theory of Evolutionary Algorithms

Organizers: Benjamin Doerr, Nikolaus Hansen, Christian Igel, and Lothar Thiele

Seminar No. 15211

Date: May 17–22, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.5.57

Participants: Youhei Akimoto, Lee Altenberg, Dirk V. Arnold, Anne Auger, Nick Barton, Roman V. Belavkin, Hans-Georg Beyer, Dimo Brockhoff, Maxim Buzdalov, Arina Buzdalova, Kenneth A. De Jong, Benjamin Doerr, Carola Doerr, Anton V. Eremeev, Carlos M. Fonseca, Tobias Friedrich, Tobias Glasmachers, Nikolaus Hansen, Hsien-Kuei Hwang, Christian Igel, Thomas Jansen, Daniel Johannsen, Joshua D. Knowles, Timo Kötzing, Oswin Krause, Marvin Künnemann, William B. Langdon, Per Kristian Lehre, Luigi Malago, Silja Meyer-Nieberg, Alberto Moraglio, Frank Neumann, Pietro S. Oliveto, Yann Ollivier, Jonathan E. Rowe, Günter Rudolph, Jonathan L. Shapiro, Peter F. Stadler, Dirk Sudholt, Lothar Thiele, Heike Trautmann, Rolf Wanka, Carsten Witt, Xin Yao, Christine Zarges

Evolutionary algorithms (EAs) are randomized search and optimization methods applicable to problems that may be non-continuous, multi-modal, noisy, multi-objective or dynamic. They have successfully been applied to a wide range of real-world applications and have demonstrated impressive performance in benchmarks for derivative-free optimization. The seminar was devoted to the theory underlying evolutionary algorithms and related methods, in order to gain a better understanding of their properties and to develop new powerful methods in a principled way. The highly international, interdisciplinary seminar brought together leading experts and young researchers in the field. The 45 participants came from 13 different countries, spread over 4 continents. Many additional researchers had expressed their interest to also attend the seminar, but could unfortunately not be considered.

Topics. The following report covers all important streams of research in the theory of evolutionary algorithms with a focus on three topics of particular current interest:

Runtime and complexity. Rigorous runtime and analysis and computational complexity theory have become the most important tools in the theory of discrete evolutionary algorithms. The Dagstuhl seminar series "Theory of Evolutionary Algorithms" has sparked this development. The drastic increase in new results, new methods, and young researchers entering this field, but also the major unsolved problems naturally lead to keeping this a focus topic.

Information geometry. Using concepts from information geometry in evolutionary algorithms is one of the most promising new theoretical direction in evolutionary computing. The seminar provided a unique opportunity to discuss perspectives and limitations of this approach.

Natural evolution. Evolutionary computing is rooted in theories of natural evolution, and many early approaches to understand basic properties of evolutionary algorithms were inspired by biological evolution theory. Still, today these two research fields are almost completely separated. We invited experts from evolution biology to help better understanding the relations between both fields. We are particularly happy that we succeeded in bringing together researchers from evolution biology and computer science in a way that was stimulating and productive.

Organization. The seminar had three types of organized presentation and discussion formats to stimulate the free discussions among the participants. There were 20–30 minutes talks on current topics followed by discussions. These included a talk on potential industrial collaborations. In addition, we had a few longer talks, which combined recent work with an overview over the state-of-the-art in a certain domain: Thomas Jansen spoke on "Understanding Randomised Search Heuristics", Nick Barton on "Limits to Adaptation", Yann Olivier introduced "Information-geometric Optimization", and Timo Kötzing presented a talk on "Stochastic Fitness Functions and Drift". Furthermore, we continued with having "breakout sessions" for longer, parallel group discussions on timely, specialized topics. These were introduced in the last seminar on "Theory of Evolutionary Algorithms". This time, these session were even more productive than previously, both because the organizers and the participants were more used to this format of interaction. The talks and breakout sessions are summarized in the full report.

We would like to thank the Dagstuhl team and the attendees for making seminar 15211 a great success and a pleasure to organize.

7.33 Multi-Disciplinary Approaches to Reasoning with Imperfect Information and Knowledge – a Synthesis and a Roadmap of Challenges

Organizers: Igor Douven, Gabriele Kern-Isberner, Markus Knauff, and Henri Prade

Seminar No. 15221

Date: May 25–29, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.5.92

Participants: Ofer Arieli, Florence Bannay-Dupin de St-Cyr, Christoph Beierle, Tarek R. Besold, Nicole Cruz de Echeverria Loebell, Igor Douven, Christian Eichhorn, Thomas Eiter, Lupita Estefania Gazzo Castaneda, Vittorio Girotto, Ulrike Hahn, Stephan Hartmann, Andreas Herzig, Anthony Hunter, Gabriele Kern-Isberner, Markus Knauff, Karolina Krzyzanowska, Henry Markovits, Odile Papini, Niki Pfeifer, Henri Prade, Marco Ragni, Hans Rott, Steven Schockaert, Henrik Singmann, Wolfgang Spohn, Sara L. Uckelman, Matthias Unterhuber, Hans Van Ditmarsch, Momme von Sydow, Emil Weydert, Gregory Wheeler

This multi-disciplinary seminar with attendees from computer science, philosophy, and psychology addressed typical problems that smart and intelligent systems in real-world scenarios have to deal with both from formal and empirical points of view. Such systems have to face, in particular, the problem of reasoning with uncertain, imprecise, incomplete, or inconsistent (in short, imperfect) information which often renders more classical, i.e., strict or deductive methods obsolete or fallacious. Reasoning with imperfect information plays a central role in practical deliberation and rational decision making. Models of human context-dependent reasoning that synthesise logical, philosophical and psychological aspects would be helpful for designing better systems. In psychology, an increasing interest in new formal methods for rational human reasoning under uncertainty can be observed, and on the other hand, philosophers and computer scientists have shown an increased attention to the experimental methods of psychology recently. In particular for computer scientists and AI researchers, it is becoming more and more interesting to see whether the systems they have been developing are materially adequate. A synthesis of rational reasoning with imperfect information that takes into account research done in artificial intelligence, but also in psychology and philosophy is needed for providing a clearer view of where we are and what are the pending issues both from computational resp. logical and cognitive viewpoints. This will help making intelligent systems more effective, and more helpful for their human users.

This seminar brought together researchers interested in rational and uncertain reasoning from a very broad scientific scope to present and discuss problems and approaches from different disciplines, consolidate common grounds, and initiate new interdisciplinary collaborations. The seminar took profit from the fact that computer scientists, philosophers, and psychologists have started quite recently to work in a common methodological paradigm with overlapping goals, converging interests, and largely shared research tools. The attendees identified challenges for new paradigms of rational reasoning, and discussed visions and foci for more interdisciplinary work.

The first day, the seminar started with (invited) survey talks on central cross-field topics, where each topic was addressed by two researchers from different disciplines:

- *Nonmonotonic reasoning and change of knowledge and beliefs*
  Marco Ragni (CS/Psy), Hans Rott (Phil)
- *Uncertain reasoning and decision theory*
  Wolfgang Spohn (Phil), Henri Prade (CS)
- *Argumentation and reasoning under inconsistency*
  Ofer Arieli (CS), Ulrike Hahn (Psy)
- *General forms of human reasoning (e.g., analogical reasoning, interpolation, and extrapolation, case-based reasoning)*
  Vittorio Girotto (Psy), Steven Schockaert (CS)

The schedule for the next days included both sessions where attendees could present and discuss their work with the audience, and time slots for discussion groups. The topics of the discussion groups were discussed in a plenary session, and four groups came out of that:

- *Topics of group 1:* Philosophers' and psychologists' view on human reasoning, and what computer scientists can contribute to that; axiomatic systems vs. psychological models – how do they fit?
- *Topics of group 2:* Empirical implications of formal reasoning systems and vice versa
- *Topics of group 3:* Combination/mixture of reasoning methods, qualitative vs. quantitative approaches; formal axiomatic systems are suitable for decision making(?)
- *Topics of group 4:* Promises and problems of probability theory; reliability, coherence, higher order probabilities

Groups 1 and 2 joined after the first session due to the closeness of the discussed topics. On Friday morning, the results of the working groups were presented, and a final, lively discussion in the plenary session closed the seminar.

Fig. 7.11
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.34 Human-Centric Development of Software Tools

Organizers: Andrew J. Ko, Shriram Krishnamurthi, Gail C. Murphy, and Janet Siegmund

Seminar No. 15222

Date: May 25–28, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.5.115

Participants: Andrew Begel, Alan Blackwell, Margaret M. Burnett, Rob DeLine, Yvonne Dittrich, Kathi Fisler, Thomas Fritz, Mark Guzdial, Stefan Hanenberg, James D. Herbsleb, Johannes Hofmeister, Reid Holmes, Christopher D. Hundhausen, Antti-Juhani Kaijanaho, Andrew J. Ko, Rainer Koschke, Shriram Krishnamurthi, Gail C. Murphy, Emerson Murphy-Hill, Brad A. Myers, Barbara Paech, Chris Parnin, Lutz Prechelt, Peter C. Rigby, Martin Robillard, Tobias Röhm, Dag Sjøberg, Andreas Stefik, Harald Störrle, Walter F. Tichy, Claes Wohlin, Thomas Zimmermann

Across our many sessions, we discussed many central issues related to research on the design of human-centric developer tools. In this summary, we discuss the key insights from each of these areas, and actionable next steps for maturing the field of human-centered developer tools.

Key Insights

Theories

Theories are a hugely important but underused aspect of our research. They help us start with an explanation, they help us explain and interpret the data we get, they help us relate our findings to others findings, and they give us vocabulary and concepts to help us organize our thinking about a phenomenon.

There are many relevant theories that we should be using:

- Attention investment is helpful in explaining why people choose to engage in programming.
- Information foraging theory helps explain where people choose to look for relevant information in code.
- Community of practice theory helps us explain how people choose to develop skills over time.

There are useful methods for generating theories, including grounded theory and participatory design. Both can result in explanations of phenomena. That said, there are often already theories about things and we don't need to engage in creating our own.

While theories are the pinnacle of knowledge, there's plenty of room for "useful knowledge" that helps us ultimately create and refine better theories. Much of the research we do now generates this useful knowledge and will eventually lead to more useful theories.

Study Recruitment

Whether developers agree to participate in a study depends on several factors:

- One factor is how much value developers perceive in participating. Value might be tangible (a gift card, a bottle of champagne), or personal (learning something from participation, or getting to share their opinion about something they are passionate about).
- Another factor in recruitment is whether the requestor is part of the developer in-group (e.g, being part of their organization, having a representative from their community conduct the research or recruit on your behalf, become part of their community before asking for their efforts)
- The cost of participating obviously has to be low, or at least low enough to account for the benefit. With these factors in mind, there are a wide range of clever and effective ways to recruit participants:

Monitor for changes in bug databases and gather data at the moment the event occurs. This makes the request timely and minimizes the cost of recall.

- Find naturalistic captures of people doing software engineering work (such as tutorials, walkthroughs, and other recorded content that developers create). This costs the nothing.
- Perform self-ethnographies or diary studies. This has some validity issues, but provides a rich source of data.
- Tag your own development work through commits to gather interesting episodes.
- Find where developers are and interview them there (e.g., the Microsoft bus stop, developer conferences), and generate low-cost, high-value ways of getting their attention (and data).

Research Questions

There was much discussion of research questions at the conference and what makes a good one. There was much agreement

that our questions should be more grounded in theories, so that we can better build upon each others' work.

Many researchers also find that the human-centered empirical studies produce results that are not directly meaningful or actionable to others. There are many possible reasons for this:

- We often don't choose research questions with more than one plausible outcome.
- We often don't report our results in a way that creates conflict and suspense. We need to show readers that there are many possible outcomes.
- We often ask "whether" questions, rather than "why" or "when" questions about tools, leading to limited, binary results, rather than richer, more subtle contributions.

Some of our research questions have validity issues that make them problematic:

- Research questions often fail to understand the populations they are asking about.
- Research questions often get involved in designing tools for people who are already designing tools for themselves. Instead, researchers should be building tools that have never existed, not building better versions of tools that already exist.

One opportunity for collaboration with researchers who are less human-centered is to collaborate on formative research that shapes the direction of research and discover new research opportunities for the field. This may create more positive perceptions of our skills, impact, and relevance to the broader fields of PL and SE.

Human-Centeredness

Historically, HCI concerns have focused on end user experiences rather than developer experiences, but HCI researchers have increasingly focused on developers and developer tools. But HCI often doesn't consider the culture and context of software engineering, and doesn't address the longitudinal / long term factors in education and skill acquisition, and so HCI may not be a sufficient lens through which to understand software engineering.

There is also a need to address low-end developers, not just "experts". Future research topics include the understand learnability of APIs, how to understand the experiences of engineers (from a sociological perspective studies such as Bucciarelli), how to think about tools from a knowledge prerequisite perspective.

Developer Knowledge Modeling

Much of what makes a developer effective is the knowledge in their mind, but we know little about what this knowledge is, how developers acquire it, how to measure and model it, and how to use these models to improve tools or enable new categories of tools. There are many open opportunities in this space that could lead to powerful new understandings about software engineering expertise and powerful new tools to support software engineering. Much of this new work can leverage research in education and learning sciences to get measures of knowledge.

Leveraging Software Development Analytics

We identified identifying different types of data that might be collected on programming processes and products. These included editing activities, compilation attempts and errors, execution attempts and errors, and check-ins. We considered ways in which these data could be enlisted to help improve teaching and learning, as well as the software development process:

- Automated interventions to improve programming processes
- Present visually to aid in decision making
- Generate notifications that could inform learners, teachers, and software developers of key events.
- Generating social recommendations.

These opportunities raise several questions:

- How do we leverage data to intervene in educational and collaborative software development settings?
- How do we design visual analytics environment to aid in decision making?
- Should interventions be automated, semi-automated, or manual? What are the trade offs?

Error Messages

We identified 5 broad classes of errors: (1) syntactic (conformance to a grammar), (2) type, (3) run-time (safety checks in a run-time system, such as array bounds, division by zero, etc.), (4) semantic (logical errors that aren't run-time errors) (5) stylistic. We distinguished between errors and more general forms of feedback, acknowledging that both needed support; in particular, each of these could leverage some common presentation guidelines.

We discussed why research has tended to focus more on errors for beginners than feedback for developers. Issues raised included the different scales of problems to diagnose across the two cases and differences in social norms around asking for help from other people (developers might be less likely to ask other people for help in order to protect their professional reputations). We discussed whether tools should report all errors or just some of them, and whether tools should try to prioritize among errors when presenting them. These had different nuances in each of students and practicing developers. We discussed the example of the coverity tool presenting only a subset of errors, since presenting all of them might lead developers to reject the tool for finding too much fault in the their code.

We discussed and articulated several principles of presenting errors: (1) use different visual patterns to distinguish different kinds of errors; (2) don't mislead users by giving incorrect advice on how to fix an error; (3) use multi-dimensional or multi-modal techniques to reveal error details incrementally; (4) when possible, allow programs to fail gently in the face of an error (example: soft typing moved type errors into run-time errors that only tripped when a concrete input triggered the error – this gives the programmer some control over when to engage with the error after it arises); (5) consider ways to allow the user to query the system to narrow down the cause of the error (rather than require them to debug the entire program).

There are several open research questions:

- Should error and feedback systems become interactive, asking the user questions to help diagnose a more concrete error (rather than report a more abstract one, as often happens with compiler syntax errors)?
- Can grammars be tailored to domain-specific knowledge to yield more descriptive error messages?
- Can patterns of variable names be used to enforce conventions and reduce the rates of some kinds of errors?
- At what point should error systems expect the user to consult with another human, rather than rely only on the computer.
- When is it more helpful to show all errors (assuming we can even compute that) versus a selection of errors? How much detail should be presented about an error at first? Does presenting all information discourage users from reading error messages?

Reviewing

Researchers in human aspects of software engineering feel a strong sense of hostility towards human-centered research, despite some recent successes in some software engineering venues. Reasons for this hostility include:

- Many human-centered researchers evaluate and critique tools without offering constructive directions forward. This creates a perception that human-centered researchers dislike or hate the research that others are doing.
- Many human-centered researchers are focused on producing understanding, whereas other researchers are focused on producing better tools. This goal mismatch causes reviewers to apply inappropriate criteria to the importance and value of research contributions.
- Many research communities in programming languages and software engineering still lack sufficient methodological expertise to properly evaluate human-centered empirical work.
- It's not explicit in reviews whether someone's methodological expertise is a good match for a paper. Expert it in a topic, not expert in a method. This leads to topic expertise matches without methodological expertise matches.
- Many challenges in reviewing come from the difference between judging a paper's validity versus judging how interesting a paper is. Non-human centered researchers do not often often find our questions interesting.

We are often our own worst enemies in reviews. We often reject each other because we're too rigid about methods (e.g., rejecting papers because of missing interrater reliability). On the other hand, we have to maintain standards. There's a lot of room for creativity in establishing rigor that is satisfying to reviewers, and we should allow for these creative ways of validating and verifying our interpretations.

Methods Training

Empirical methods are not popular to learn. However, when our students and colleagues decide to learn them, there are many papers, textbooks, classes and workshops for learning some basic concepts in human-subjects software engineering research.

There are many strategies we might employ to broadly increase methodological expertise in our research communities:

- We should spend more time in workshops and conferences teaching each other how to do methods well.
- Software engineers need to learn empirical methods too, and teaching them as undergraduates will lead to increased literacy in graduate students.
- There is much we can do to consolidate and share teaching resources that would make this instruction much more efficient.
- HCI research methods are broadly applicable and there are many more places to learn them.

There aren't good methods for researching learning issues yet. Moreover, most of these methods cannot be learned quickly. We must devise ways of teaching these methods to students and researchers over long periods of time.

Fig. 7.12
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.35 Computational Social Choice: Theory and Applications

Organizers: Britta Dorn, Nicolas Maudet, and Vincent Merlin

Seminar No. 15241

Date: June 7–12, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.6.1

Participants: Haris Aziz, Dorothea Baumeister, Péter Biró, Sylvain Bouveret, Steven J. Brams, Felix Brandt, Robert Bredereck, Markus Brill, Ioannis Caragiannis, Katarína Cechlárová, Jiehua Chen, Yann Chevaleyre, Andreas Darmann, Britta Dorn, Edith Elkind, Ulle Endriss, Piotr Faliszewski, Christian Geist, Judy Goldsmith, Umberto Grandi, Paul Harrenstein, Christian Klamler, Bettina Klaus, Flip Klijn, Dominikus Krüger, Jérôme Lang, Annick Laruelle, Michel Le Breton, David Manlove, Nicholas Mattei, Nicolas Maudet, Vincent Merlin, Trung Thanh Nguyen, Francesca Rossi, Jörg Rothe, Ildikó Schlotter, Arkadii Slinko, Toby Walsh, Gerhard J. Woeginger, Jia Yuan Yu, William S. Zwicker

Computational social choice is an interdisciplinary research area dealing with the aggregation of preferences of groups of agents in order to reach a consensus decision that realizes some social objective. Economists typically view markets as an optimal mean for coordinating the activities and allocation of resources across a group of heterogeneous agents based on their utilities or preferences. By contrast, the methods of social choice, broadly defined, focus on coordination mechanisms that do not rely on prices, monetary/resource transfer or market structures, while still defining social objectives that account for individual preferences. Some classic (but certainly not exhaustive) topics of study in social choice topics include:

- voting procedures, where a single alternative must be taken given the preferences of individuals group members;
- fair division, which deals with the distribution of goods among a group reflecting individual preferences and fairness criteria;
- matching problems, in which agents/items are matched in a way that respects both preferences and other constraints.

The theoretical treatment of these problems is concerned with the existence of solutions which could be defended on normative grounds. In classical social choice, desirable solution concepts satisfy certain properties, such as: efficiency, non-discriminatory treatment of agents; envy-freeness; stability (or equilibrium) with respect to incentives; non-manipulability; and a variety of others. Over the past 15 years, the computational properties of these solution concepts have emerged as critically important to their theoretical viability and practical impact. Computer scientists in both AI and theoretical computer science have developed efficient algorithms for realizing certain social choice functions, proven the computational intractability of others, studied the theoretical and practical communication requirements of these procedures, and developed computational tools to sharpen our understanding of incentives, manipulation, and other important phenomena. Applying a computational lens to these theoretical investigations has led to breakthroughs that have supported a variety of real-world applications like web-page ranking, fair buy-sell/exchange protocols, and the development of much more socially efficient exchanges for organ transplantation.

At the same time, the era of networked communication and "big data" has made it easier than ever to infer people's preferences and have them engage with ever larger groups. This has opened up tremendous opportunities for the application of social choice to a wider range of "lower stakes, higher frequency" group decisions. Hence, it introduces new challenges for social choice – many mechanisms for the problems above have been designed using assumptions that – while suitable for "high stakes" domains like political voting, or matching in labor markets and organ donations – are entirely untenable in other domains.

The objective of the seminar was to continue the series of meetings on theoretical computational social choice previously held in Dagstuhl, but the emphasis was on problems which have practical relevance. We have addressed in particular three lines of works concerning issues in social choice: voting, matching, and fair division. The seminar brought together 41 researchers from 18 countries and various fields such as computer science, mathematics, social choice theory, economics, political sciences, and industry. The meeting gathered both participants focusing on the theoretical foundations of computational social choice, and those seeking to apply social choice mechanisms to real-world problems of both the high-stakes/low-frequency and the low-stakes/high-frequency variety.

The technical program of the seminar included overview talks, regular seminar talks, a rump session and slots for communication and work on open problems. The three overview talks presented open questions and challenges in multiwinner voting (complemented by a panel discussion), economics and computa-

tion, and in matching in the context of assignments of teachers to schools. The 26 regular seminar talks covered the three lines of work concerning voting, matching, and fair division/resource allocation. Current trends in these fields as reflected by the contributions include allocation of indivisible items under ordinal preferences, the study of well-behaved preference structures (e.g. single-peaked, single-crossing), multiwinner elections, mixed voting systems, and several highly challenging special cases of matching problems. Challenges from real-world applications included online fair division for the distribution of food donations to a food bank, assignments of referees to papers for scientific reviewing, peer grading in massive online open courses, online voting and online participation, sharing cars, junior doctor allocation, and house swapping. Furthermore, several online platforms dedicated to social choice were presented and discussed during the seminar. Precious feedback was collected by the teams of developers.

The program offered the possibility to present open problems and provided slots for working groups on these topics as well as a final session for presentation of outcomes. Several working groups were formed some of which obtained first results during the seminar week. The research projects initiated in these groups are still ongoing. Many participants also used these slots for collaboration with their co-authors that were present at the seminar.

The rump session consisted of 17 five minute contributions, ranging from announcements of events related to the community, over presentation of tools for preference aggregation and online voting, preference libraries for datasets, applications like sharing cars, to short research talks and presentations of open problems.

To conclude, the seminar acknowledged that more and more contributions in computational social choice are driven by real world issues, with many potential applications for industry and policy making. It confirmed that theoretical considerations enable, justify and guarantee the quality of practical applications. Conversely, the specific features and constraints of applications provide novel theoretical challenges and new directions for foundational research.

The participants greatly appreciated the time devoted to working group sessions and benefited from the seminar in various ways: by learning about new problems, many of them being directly inspired from real world issues; by being introduced to several existing tools; by having the possibility to interact and to develop new collaborations. A next event will be the COMSOC workshop in Toulouse in June 2016. It will be co-located with the meeting of the COST Action IC1205 “Computational Social Choice”, and one day will be devoted to applications and interactions with industry, in line with the 15241 Dagstuhl seminar.

We would like to thank all participants for their contributions, discussions, ideas and collaborations, making this seminar a very productive and enjoyable one. In particular, we sincerely thank the team of Schloss Dagstuhl for the great support and excellent organization.

7.36 Complexity of Symbolic and Numerical Problems

Organizers: Peter Bürgisser, Felipe Cucker, Marek Karpinski, and Nicolai Vorobjov

Seminar No. 15242

Date: June 7–12, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.6.28

Participants: Dennis Amelunxen, Saugata Basu, Peter Bürgisser, Michel Coste, Felipe Cucker, James H. Davenport, Dima Grigoriev, Edward A. Hirsch, Erich Kaltofen, Marek Karpinski, Pascal Koiran, Martin Lotz, Klaus Meer, Friedhelm Meyer auf der Heide, Dmitrii V. Pasechnik, Vladimir Podolskii, Ilia Ponomarenko, Natacha Portier, Marie-Françoise Roy, Mohab Safey El Din, Éric Schost, Vladimir Shpilrain, Anatol Slissenko, Thorsten Theobald, Nicolai Vorobjov, Andreas Weber

The seminar was dedicated to Prof. Dima Grigoriev on the occasion of his 60th birthday. Its aim was to discuss modern trends in computational real algebraic geometry, in particular, areas related to solving real algebraic and analytic equations and inequalities. Very recent new developments in the analysis of these questions from the point of view of *tropical mathematics* were also presented.

Historically there were two strands in the computational approach to polynomial systems' solving. One is the tradition of numerical analysis, a classical achievement of which is the *Newton's method*. Various other approximation algorithms were developed since then, some based on the idea of a *homotopy*. Numerical analysis did not bother to introduce formal models of computations (and hence computational complexity considerations) but developed refined methods of estimations of convergency rates. Another tradition emerged from algebra, particularly in classical works of Cayley, Sylvester and Macaulay. Algebraic results concerning *real* solutions go further back to the Descartes' rule and Sturm sequences. An important contribution to the subject from logic was Tarski's constructive quantifier elimination procedures for algebraically closed and real closed fields. The computations considered in this tradition are exact, under modern terminology – "symbolic". They naturally fit into standard models of computation (Turing Machines, straight-line programs, computation trees) thus lending themselves to complexity analysis.

Until 1990s these two strands developed largely independently. One of the important unifying ideas became the concept of a *real numbers* (or *BSS*) machine suggested by Blum, Shub and Smale which can be considered as a model of computation for the numerical analysis. This idea led to Smale's 9th and 17th problems, which became an inspiration for many researchers in the field.

The seminar considered a wide set of questions related to the current state of the symbolic and numeric approaches to algorithmic problems of real algebraic and analytic geometry, also from the novel perspective of tropical and max/plus mathematics.

Fig. 7.13
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.37 Sparse Modelling and Multi-Exponential Analysis

Organizers: Annie Cuyt, George Labahn, Avraham Sidi, and Wen-shin Lee

Seminar No. 15251

Date: June 14–19, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.48

Participants: John Abbott, Fredrik Andersson, Andrew Arnold, Dmitry Batenkov, Bernhard Beckermann, Matteo Briani, Adhemar Bultheel, Mathieu Collowald, Annie Cuyt, Stefano de Marchi, Pier Luigi Dragotti, Jürgen Gerhard, Mark Giesbrecht, Karlheinz Gröchenig, Lasse Holmström, Evelyne Hubert, Erich Kaltofen, Stefan Kunis, George Labahn, Wen-Shin Lee, David Levin, David Li, Daniel Lichtblau, Ivan Markovsky, Ana C. Matos, Katharine M. Mullen, Luca Perotti, Thomas Peter, Gerlind Plonka-Hoch, Daniel Potts, Daniel Roche, Avraham Sidi, Manfred Tasche, Akira Terui, Konstantin Usevich, Ulrich von der Ohe, Katrin Wannenwetsch, Stephen M. Watt, Marius Wischerhoff, Yosef Yomdin, Lihong Zhi

The seminar brought together a number of researchers from polynomial interpolation, rational approximation and exponential analysis. The five day seminar centered around talks on Exponential Analysis (Day 2), Rational Approximation (Day 3) and Sparse Interpolation (Day 4). Applications were grouped on Day 1 in order to challenge the participants to discuss them further while related topics, mainly from Numerical Linear Algebra, were scheduled on Day 5.

The seminar itself started with a talk by Cuyt and Lee pointing out the considerable intersection of the three main themes, particularly as they all strongly overlap. In order to reach out to industry and connect the scientific research to the industrial needs, several participants working at industrial or real-life applications were invited for a presentation on the first day of the seminar. Then interaction about these topics would occur naturally throughout the week. We mention talks on Mobile sampling and sensor networks (Karlheinz Gröchenig), High-speed fluorescence lifetime imaging (David Li), The estimation of variable star periods (Daniel Lichtblau) and Imaging of structured arrays (Adhemar Bultheel).

In the past the three communities have mostly been following distinct paths of research and methods for computation. One of the highlights of the seminar was the realization of significant commonalities between the communities, something nicely pointed out in the talk of Roche. Prony's method takes center stage in this case, with its origin in 1795 being used to solve problems in exponential analysis. Prony's method appeared much later in the case of sparse polynomial interpolation with it's use by Blahout, Ben-or/Tiwari, and Giesbrecht/Labahn/Lee. Prony's method takes samples at multiples of a common point to determine the support and then makes use of separate Hankel methods for determining the individual coefficients or weights of the expression.

Numerical conditioning was a significant issue in many talks at the seminar. Beckermann and later Matos looked at numerical conditioning of Padé and rational approximation problems. In the former case Beckermann used the close relationship of Padé approximation to Prony's method to point out that the latter is, for the most part, a provably ill-conditioned problem. Still there were a number of approaches in both areas which attempted to address this conditioning issue. In the case of numerical computation of sparse polynomial interpolants, use is made of randomization to produce a better conditioning of the problem, primarily by separating the roots appearing in Prony's problem. A similar idea also appears in exponential analysis making use of the notion of stride length. In both cases the object is to spread out the roots which arise in Prony's method.

Rather than spreading out the roots one can instead spread out the coefficients of a sparse polynomial/exponential expression for improving numerical performance. Sparse interpolation does this by making use of the concept of diversification where the coefficients are spread out multiplying evaluation points using a random multiplier. A corresponding concept in exponential analysis is the use of shifted samples which is useful to address the problem of anti-aliasing.

Sparse interpolation also makes use of the concept of small primes sparse interpolation where exponents are reduced modulo a small prime. This recovers the exponents modulo the small prime. Doing this for a number of small primes (which can be done in parallel) allows one to reconstruct the true exponents. Of course one encounters the problem of collisions and inadvertant combinations of exponents. It was noticed at the seminar that exponential analysis has a corresponding technique which made use of sub-sampling. Collisions in this case correspond to aliasing. Again the different communities reported on their methods for overcoming such collisions/aliasing problems.

Researchers at the seminar also showed interest in multivariate Prony methods. In the case of sparse interpolation one

encounters Zippel's method while in exponential analysis there are projection methods. In these cases one attempts recursive methods for estimating the support of the underlying multivariable expression. In the case of multivariate polynomial interpolation a second approach is to convert the multivariate problem into a univariate problem by making use of randomized Kronecker substitution. Exponential sums takes a similar approach using random lattice projection.

While there were strong commonalities between the main research areas, there were also some strong differences between the topics noted at the seminar. The most telling of these differences was the analysis of exponential sums which have polynomial, rather than constant coefficients. Such expressions appear naturally when modeling solutions of linear differential equations where the associated polynomial has repeated roots. Of course such problems have considerable numerical issues when the roots of the associated polynomial are close but not numerically equal. Sidi and Batenkov both pointed out the importance and difficulties when dealing with such problems.

The seminar was also important for illustrating the applications of the three research areas. In many cases the applications involved the need to only work with sums having a small sparse support rather than with the complete set of possible nonzero elements. Methods from the multivariate Prony problem were exploited by Collowald and Hubert to determine new cubature formulas invariant to some specific finite groups action. Markovsky showed the similarities to the exponential sum problems with the notion of low rank approximation of structured matrices. Software was also discussed. Numerical analysis of errors on experimental runs also brought up the issue of the type of random distributions used when simulating errors for the experiments.

7.38 Logics for Dependence and Independence

Organizers: Erich Grädel, Juha Kontinen, Jouko Väänänen, and Heribert Vollmer

Seminar No. 15261

Date: June 21–26, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.6.6.70

Participants: Faried Abu Zaid, Dietmar Berwanger, Olaf Beyersdorff, Julian Bradfield, Maurice Chandoo, Ivano Alessandro Ciardelli, Anuj Dawar, Nicolas de Rugy-Altherre, Arnaud Durand, Fredrik Engström, Valentin Goranko, Erich Grädel, Miika Hannula, Lauri Hella, Åsa Hirvonen, Antti Hyttinen, Phokion G. Kolaitis, Juha Kontinen, Antti Kuusisto, Sebastian Link, Martin Lück, Kerkko Luosto, Allen L. Mann, Arne Meier, Martin Otto, Gianluca Paolini, Raine Rönnholm, Katsuhiko Sano, Svenja Schalthöfer, Thomas Schwentick, Ilya Shpitser, Alex Simpson, Johanna Stumpf, Bernhard Thalheim, Jouko Väänänen, Jonni Virtema, Heribert Vollmer, Fan Yang

Brief Introduction to the Topic

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

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

The Dagstuhl seminar "Dependence Logic: Theory and Applications" had a major impact to the field of dependence logic opening up connections to new application areas. The aim of this follow-up seminar was to gather together the people working in dependence logic and in the application areas, especially those researchers who have recently started working in this quickly developing area to communicate state-of-the-art advances and embark on a systematic interaction.

Organization of the Seminar and Activities

The seminar brought together 38 researchers from mathematics, statistics, database theory, natural language semantics, and theoretical computer science. The participants consisted of both senior and junior researchers, including a number of postdocs and advanced graduate students.

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

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

- Jouko Väänänen and Juha Kontinen, Dependence Logic
- Bernhard Thalheim, Database Constraints – A Survey
- Ilya Shpitser, Causal inference
- Lauri Hella, Modal dependence logic
- Ivanio Ciardelli, Dependency as Question Entailment
- Antti Hyttinen, Statistical Independence, Causality and Constraint Satisfaction

There were additionally two introductory talks with a more focused and technical topic:

- Alex Simpson, Sheaf semantics for independence logics
- Phokion Kolaitis, The Query Containment Problem: Set Semantics vs. Bag Semantics

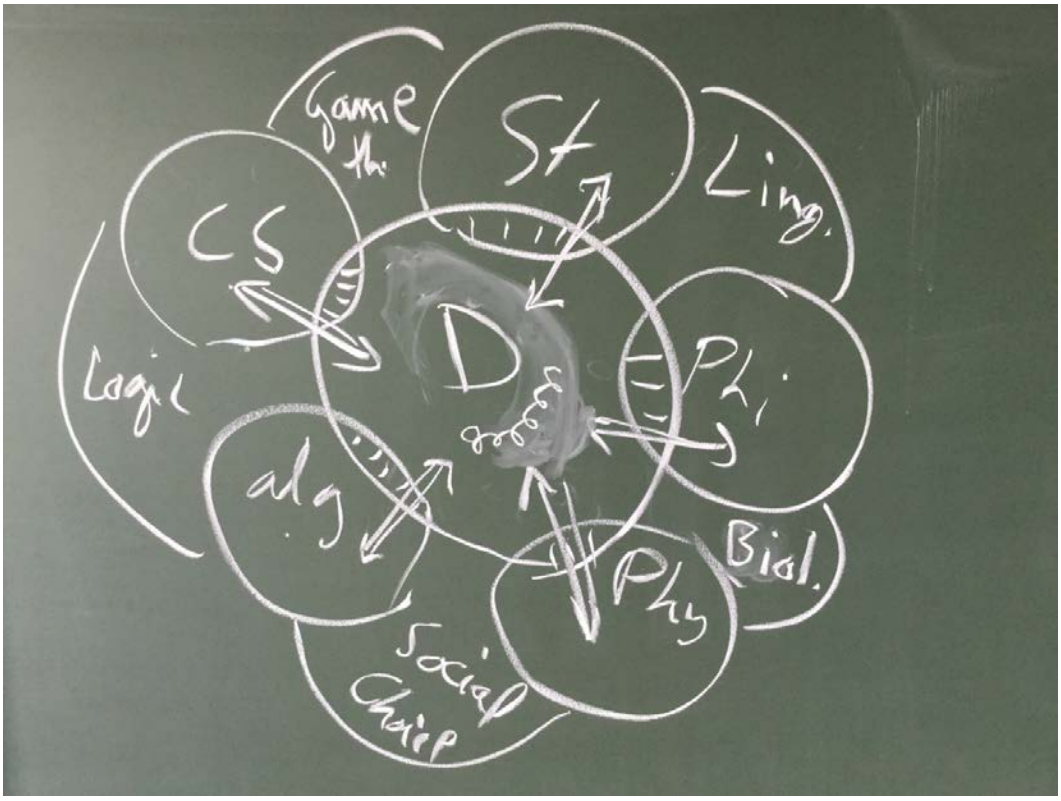

Fig. 7.14
The blackboard after Jouko Väänänen's conclusion of the seminar.

Additionally, the following shorter presentations were given during the seminar:

- Åsa Hirvonen, Model theoretic independence
- Kerkko Luosto, Dimensions for Modal Dependence Logic
- Gianluca Paolini, Measure teams
- Olaf Beyersdorff, Proof Complexity of Quantified Boolean Formulas
- Antti Kuusisto, Propositional dependence logic via Kripke semantics
- Johanna Stumpf, Characterisation of the expressive power of modal logic with inclusion atoms
- Sebastian Link, Dependence-driven, non-invasive cleaning of uncertain data
- Jonni Virtema, Complexity of Propositional Inclusion and Independence Logic
- Katsuhiko Sano, Characterizing Frame Definability in Team Semantics via The Universal Modality
- Raine Rönnholm, Expressing properties of teams in k-ary inclusion-exclusion logic
- Julian Bradfield, On the structure of events in Boolean games
- Fan Yang, Some proof theoretical results on propositional logics of dependence and independence
- Erich Grädel, Counting in Team Semantics
- Fredrik Engström, Generalized quantifiers and Dependence Logic
- Miika Hannula, Axiomatizing dependencies in team semantics
- Dietmar Berwanger, An NL-fragment of inclusion logic
- Nicolas de Rugy-Altherre, Tractability Frontier of Data Complexity in Team Semantics

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

Concluding Remarks and Future Plans

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

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

7.39 Life-Long Health Behavior-Change Technologies

Organizers: Susanne Boll, Eric Hekler, and Predrag Klasnja
Seminar No. 15262

Date: June 21–26, 2015 | Dagstuhl Seminar

Participants: Audie Atienza, Jakob E. Bardram, Susanne Boll, Parisa Eslambolchilar, Eric Hekler, Hermie Hermens, Wilko Heuten, Katherine Kim, Andrew Krause, Kai Kunze, Cesar Martin-Moreno, Jochen Meyer, Florian "Floyd" Mueller, Sean Munson, Harri Oinas-Kukkonen, Misha Pavel, Alison Phillips, Harald Reiterer, Britta Renner, Jochen Schnauber, m.c. schraefel, Ralf Schwarzer, Saul Shiffman, Katie A. Siek, Donna Spruijt-Metz, Ambuj Tewari, Andi Winterboer

(No documentation available)

Fig. 7.15
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of (above) TreeState Productions GmbH, www.treestate.de and (below) Dr. Christian Lindig.

7.40 Algorithms and Scheduling Techniques to Manage Resilience and Power Consumption in Distributed Systems

Organizers: Henri Casanova, Ewa Deelman, Yves Robert, and Uwe Schwiegelshohn

Seminar No. 15281

Date: July 5–10, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.7.1

Participants: Sadaf Alam, Olivier Beaumont, Anne Benoit, George Bosilca, Henri Casanova, Ewa Deelman, Rafael Ferreira da Silva, Carsten Franke, Michael Gerhards, Dominik Göddeke, Julien Herrmann, Sascha Hunold, Fredy Juarez, Kate Keahey, Thilo Kielmann, Dieter Kranzlmüller, Julien Langou, Maciej Malawski, Loris Marchal, Satoshi Matsuoka, Rami Melhem, Bernd Mohr, Alix Munier, Jaroslaw Nabrzyski, Wolfgang E. Nagel, vManish Parashar, Ilia Pietri, Sabri Pllana, Suraj Prabhakaran, Padma Raghavan, Yves Robert, Stephan Schlagkamp, Uwe Schwiegelshohn, Oliver Sinnen, Renata Slota, Veronika Sonigo, Leonel Sousa, Andrei Tchernykh, Jeffrey S. Vetter, Frédéric Vivien, Felix Wolf

Many computer applications are executed on large-scale systems that comprise many hardware components, such as clusters that can be federated into distributed cloud computing or grid computing platforms. The owners/managers of these systems face two main challenges: failure management and energy management.

Failure management, the goal of which is to achieve resilience, is necessary because a large number of hardware resources implies a large number of failures during the execution of an application. While hardware resources can be made more reliable via the use of redundancy, this redundancy increases component cost. As a result, systems deployed within budget constraints must be built from unreliable components, that have a finite Mean Time Between Failure (MTBF), i.e., commercial-of-the-shelf components. For instance, a failure would occur every 50 minutes in a system with one million components, even if the MTBF of a single component is as large as 100 years.

Energy management, the goal of which is to optimize power consumption and to handle thermal issues, is also necessary due to both monetary and environmental constraints. While in today's systems, processors are the most power-consuming components, it is anticipated that in future distributed systems, the power dissipated to perform communications and I/O transfers will make up a much larger share of the overall energy consumption. In fact, the relative cost of communication is expected to increase dramatically, both in terms of latency/overhead and of consumed energy. Consequently, the computation and communication workloads of typical applications executed in HPC and/or cloud environments will lead to large power consumption and heat dissipation.

These two challenges, resilience and energy efficiency, are currently being studied by many researchers. Some of these researchers come from a "systems" culture, and investigate in particular systems design and management strategies that enhance resilience and energy efficiency. These strategies include high-level resource-provisioning policies, pragmatic resource allocation and scheduling heuristics, novel approaches for designing and deploying systems software infrastructures, and tools for monitoring/measuring the state of the system. Other researchers come from an "algorithms" culture. They investigate formal definitions of resilience and energy efficiency problems, relying on system models of various degrees of accuracy and sophistication, and aiming to obtain strong complexity results and algorithmic solutions for solving these problems. These two communities are quite often separated in the scientific literature and in the field. Some of the pragmatic solutions developed in the former community appear algorithmically weak to the latter community, while some of the algorithmic solutions developed by the latter community appear impractical to the former community. Furthermore, the separation of application and system platform due to ubiquitous resource virtualization layers also interferes with an effective cooperation of algorithmic and system management methods, and in particular to handle resiliency and energy efficiency. To move forward, more interaction and collaboration is needed between the systems and the algorithms communities, an observation that was made very clear during the discussions in the predecessor Dagstuhl seminar [45].

The broader challenge faced by systems and algorithms designer is that the optimization metrics of interest (resilience, power consumption, heat distribution, performance) are intimately related. For instance, high volatility in power consumption due to the use of dynamic frequency and voltage scaling (DFVS)

[45] Dagstuhl Seminar 13381, *Algorithms and Scheduling Techniques for Exascale Systems* (2013), organized by Henri Casanova, Yves Robert and Uwe Schwiegelshohn (http://www.dagstuhl.de/13381).

is known to lead to thermal hotspots in a datacenter. Therefore, the datacenter must increase the safety margin for their cooling system to handle these hotspots. As a result, the power consumed by the cooling system is increased, possibly increasing the overall power consumption of the whole system, even though the motivation for using DVFS in the first place was to reduce power consumption! When resilience is thrown into the mix, then the trade-offs between the conflicting resilience, performance, and energy goals become even more intertwined. Adding fault-tolerance to a system, for instance, by using redundant computation or by periodically saving the state of the system to secondary storage, can decrease performance and almost always increases hardware resource requirements and thus power consumption. The field is rife with such conundrums, which must be addressed via systems and algorithms techniques used in conjunction. In this seminar, we have brought together researchers and practitioners from both the systems and the algorithms community, so as to foster fruitful discussions of these conundrums, many of which were touched upon in the predecessor seminar but by no means resolved.

To provide a clear context, the seminar focused around workflow applications. Workflows correspond to a broad and popular model of computation in which diverse computation tasks (which many themselves follow arbitrary models of computation) are interconnected via control and data dependencies. They have become very popular in many domains, ranging from scientific to datacenter applications, and share similar sets of challenges and current solutions. Part of the motivation of using workflows, and thus to develop workflow management systems and algorithms, is that they make it possible to describe complex and large computations succinctly and portably. Most of the invited seminar participants have worked and are currently working on issues related to the efficient, resilient, and energy efficient execution of workflows in distributed platforms. They thus provide an ideal focal and unifying theme for the seminar.

A number of workflow tools is available to aid the users in defining and executing workflow applications. While these tools are thus designed primarily to support the end user, they are in fact ideal proving grounds for implementing novel systems and algorithms techniques to aim at optimizing performance, resilience, and energy efficiency. Therefore, these tools provide a great opportunity to enhance both the application and the software infrastructure to meet both the needs of the end users and of the systems owners/managers. These goals are very diverse and, as we have seen above, intertwined, so that re-designing algorithms and systems to meet these goals is a difficult proposition (again, higher resilience often calls for redundant computations and/or redundant communication, which in turn consumes extra power and can reduce performance). In a broad sense, we are facing complex multi-criteria optimization problems that must be (i) formalized in a way that is cognizant of the practical systems constraints and hardware considerations; (ii) solved by novel algorithms that are both fast (so that they can be used in an on-line manner) and robust (so that they can tolerated wide ranges of scenarios with possibly inaccurate information).

The goal of this seminar was to foster discussions on, and articulate novel and promising directions for addressing the challenges highlighted above. International experts in the field have investigated how to approach (and hopefully at least partially address) the challenges that algorithms and system designers face due to frequent failures and energy usage constraints. More specifically, the seminar has addressed the following topics:

- Multi-criteria optimization problems as applicable to fault-tolerance / energy management
- Resilience techniques for HPC and cloud systems
- Robust and energy-aware distributed algorithms for resource scheduling and allocation in large distributed systems.
- Application-specific approaches for fault-tolerance and energy management, with a focus on workflow-based applications

Although the presentations at the seminar were very diverse in scope, ranging from practice to theory, an interesting observation is that many works do establish strong links between practice (e.g., particular applications, programming models) and theory (e.g., abstract scheduling problems and results). In particular, it was found that workflow applications, far from being well-understood, in fact give rise to a range of interrelated and interesting practical and theoretical problems that must be solved conjointly to achieve efficiency at large scale. Estimating task weights, scheduling with uncertainties, mapping at scale, remapping after failures, trading performance and energy, these are a few challenges that have been discussed at length during the seminar. Such observations make it plain that forums that blends practice and theory, as is the case with this seminar, are very much needed.

The seminar brought together 41 researchers from Austria, France, Germany, Japan, Netherlands, New Zealand, Poland, Portugal, Spain, Sweden, Switzerland, UK and USA, with interests and expertise in different aspect of parallel and distributed computing. Among participants there was a good mix of senior researchers, junior researchers, postdoctoral researchers, and Ph.D. students. Altogether there were 29 presentations over the 5 days of the seminar, organized in morning and late-afternoon sessions. The program was as usual a compromise between allowing sufficient time for participants to present their work, while also providing unstructured periods that were used by participants to pursue ongoing collaborations as well as to foster new ones. The feedback provided by the participants show that the goals of the seminar, namely to circulate new ideas and create new collaborations, were met to a large extent.

The organizers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and support in the arrangement of a very successful and productive event.

7.41 The Constraint Satisfaction Problem: Complexity and Approximability

Organizers: Andrei A. Bulatov, Venkatesan Guruswami, Andrei Krokhin, and Dániel Marx

Seminar No. 15301

Date: July 19–24, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.7.22

Participants: Albert Atserias, Per Austrin, Libor Barto, Manuel Bodirsky, Jonah Brown-Cohen, Andrei A. Bulatov, Catarina Alexandra Carvalho, Siu On Chan, Hubie Chen, Victor Dalmau, Laszlo Egri, Serge Gaspers, Leslie Ann Goldberg, Venkatesan Guruswami, Mark R. Jerrum, Peter Jonsson, Alexandr Kazda, Vladimir Kolmogorov, Marcin Kozik, Andrei Krokhin, Euiwoong Lee, Konstantin Makarychev, Yury Makarychev, Rajsekar Manokaran, Barnaby Martin, Dániel Marx, Neeldhara Misra, Joanna Ochremiak, Michael Pinsker, David Richerby, Michal Rolinek, Francesco Scarcello, David Steurer, Stefan Szeider, Johan Thapper, Madhur Tulsiani, Matt Valeriote, Magnus Wahlström, Ross Willard, Yuichi Yoshida, Dmitriy Zhuk, Stanislav Zivny

The *constraint satisfaction problem*, or CSP in short, provides a unifying framework in which it is possible to express, in a natural way, a wide variety of computational problems dealing with mappings and assignments, including satisfiability, graph colorability, and systems of equations. The CSP framework originated 25–30 years ago independently in artificial intelligence, database theory, and graph theory, under three different guises, and it was realised only in the late 1990s that these are in fact different faces of the same fundamental problem. Nowadays, the CSP is extensively used in theoretical computer science, being a mathematical object with very rich structure that provides an excellent laboratory both for classification methods and for algorithmic techniques, while in AI and more applied areas of computer science this framework is widely regarded as a versatile and efficient way of modelling and solving a variety of real-world problems, such as planning and scheduling, software verification and natural language comprehension, to name just a few. An instance of CSP consists of a set of variables, a set of values for the variables, and a set of constraints that restrict the combinations of values that certain subsets of variables may take. Given such an instance, the possible questions include (a) deciding whether there is an assignment of values to the variables so that every constraint is satisfied, or optimising such assignments in various ways, (b) counting satisfying assignments, exactly or approximately, or (c) finding an assignment satisfying as many constraints as possible. There are many important modifications and extensions of this basic framework, e.g. those that deal with valued or global constraints.

Constraint satisfaction has always played a central role in computational complexity theory; appropriate versions of CSPs are classical complete problems for most standard complexity classes. CSPs constitute a very rich and yet sufficiently manageable class of problems to give a good perspective on general computational phenomena. For instance, they help to understand which mathematical properties make a computational problem tractable (in a wide sense, e.g. polynomial-time solvable or non-trivially approximable, fixed-parameter tractable or definable in a weak logic). It is only natural that CSPs play a role in many high-profile conjectures in complexity theory, exemplified by the Dichotomy Conjecture of Feder and Vardi and the Unique Games Conjecture of Khot.

The recent flurry of activity on the topic of the seminar is witnessed by three previous Dagstuhl seminars, titled "Complexity of constraints" (06401) and "The CSP: complexity and approximability" (09441, 12541), that were held in 2006, 2009, and 2012 respectively. This seminar was a follow-up to the 2009 and 2012 seminars. Indeed, the exchange of ideas at the 2009 and 2012 seminars has led to new ambitious research projects and to establishing regular communications channels, and there is a clear potential of a further systematic interaction that will keep on cross-fertilizing the areas and opening new research directions. The 2015 seminar brought together forty three researchers from different highly advanced areas of constraint satisfaction and involved many specialists who use universal-algebraic, combinatorial, geometric and probabilistic techniques to study CSP-related algorithmic problems. The participants presented, in 28 talks, their recent results on a number of important questions concerning the topic of the seminar. One particular feature of this seminar is a significant increase in the number of talks involving multiple subareas and approaches within its research direction – a definite sign of the growing synergy, which is one of the main goals of this series of seminars.

Concluding Remarks and Future Plans. The seminar was well received as witnessed by the high rate of accepted invitations and the great degree of involvement by the participants. Because of the multitude of impressive results reported during the seminar and the active discussions between researchers with different expertise areas, the organisers regard this seminar as

a great success. With steadily increasing interactions between such researchers, we foresee a new seminar focussing on the interplay between different approaches to studying the complexity and approximability of the CSP. Finally, the organisers wish to express their gratitude to the Scientific Directors of the Dagstuhl Centre for their support of the seminar.

Description of the Topics of the Seminar

Classical computational complexity of CSPs. Despite the provable existence of intermediate (say, between P and NP-complete, assuming P $\neq$ NP) problems, research in computational complexity has produced a widely known informal thesis that "natural problems are almost always complete for standard complexity classes". CSPs have been actively used to support and refine this thesis. More precisely, several restricted forms of CSP have been investigated in depth. One of the main types of restrictions is the *constraint language* restriction, i.e., a restriction on the available types of constraints. By choosing an appropriate constraint language, one can obtain many well-known computational problems from graph theory, logic, and algebra. The study of the constraint language restriction is driven by the CSP *Dichotomy Conjecture* of Feder and Vardi which states that, for each fixed constraint language, the corresponding CSP is either in P or NP-complete. There are similar dichotomy conjectures concerning other complexity classes (e.g. L and NL). Recent breakthroughs in the complexity of CSP have been made possible by the introduction of the universal-algebraic approach, which extracts algebraic structure from the constraint language and uses it to analyse problem instances. The above conjectures have algebraic versions which also predict in algebraic terms where the boundary between harder problems and easier problems lies. The algebraic approach has been applied to prove the Dichotomy Conjecture in many important special cases (e.g. Bulatov's dichotomy theorems for 3-valued and conservative CSPs), but the general problem remains open. Barto and Willard described the current state-of-the-art in proving this conjecture, gave insights into the main stumbling blocks (notably, the convoluted ways in which systems of linear equations appear in constraint problems), and outlined avenues of attack on those obstacles. Kozik gave a new simplified algorithm for CSPs solvable by local consistency methods, confirming an earlier conjecture. Brown-Cohen presented new results leading to closer interchange of ideas between algebraic and probabilistic approaches to CSPs.

Valued CSP is a significant generalisation of CSP that involves both feasibility and optimisation aspects. The complexity of language-based restriction for VCSPs was considered in the talks by Kolmogorov, Thapper, and Živný. Very strong result in this direction were reported, especially the full description of tractable cases modulo CSP, which closes a sequence of strong and unexpected results on VCSPs obtained during last five years.

The complexity of counting solutions for CSPs, with many results, was investigated by Goldberg, Jerrum, and Richerby.

Along with the constraint language restriction on CSP, the other main type is the structural restriction (i.e. restriction on the immediate interaction between variables in instances). Structural restrictions leading to tractability are well-understood, by results of Grohe and Marx. The so-called "hybrid" tractability in CSP, which is tractability that cannot be attributed to a constraint language restriction or to a structural restriction alone, has not received a great deal of attention yet, and is one of the possible avenues of future work. Rolínek, Scarcello, and Živný described recent results on hybrid tractability for CSPs and VCSPs, including counting problems.

Approximability of CSPs. The use of approximation algorithms is one of the most fruitful approaches to coping with NP-hardness. Hard optimization problems, however, exhibit different behavior with respect to approximability, making it an exciting, and by now, well-developed but far from fully understood, research area. The CSP has always played an important role in the study of approximability. For example, it is well known that the famous PCP theorem has an equivalent reformulation in terms of inapproximability of a certain CSP; moreover, the recent combinatorial proof of this theorem by Dinur in 2006 deals entirely with CSPs. The first optimal inapproximability results by Håstad in 2001 were about certain CSPs, and they led to the study of a new hardness notion called *approximation resistance* (which, intuitively, means that a problem cannot be approximated beyond the approximation ratio given by picking an assignment uniformly at random, even on almost satisfiable instances). Many CSPs have been classified as to whether they are approximation resistant but there is not even a reasonable conjecture for a full classification. Lee and Tulsiani presented new results on approximation resistance.

Many approximation algorithms for CSPs are based on the Sum-of-Squares method, Linear Programming and Semidefinite Programming. Recent developments in proving lower bounds for such algorithms were presented by Chan and Steurer.

Improved approximation algorithms for certain infinite-domain CSPs related to correlation clustering were given by K. Makarychev.

New applications of algebraic approach to investigate approximability of CSPs were given by Austrin and Dalmau.

Parameterized complexity of CSPs. A different way to cope with NP-hardness is provided by parameterized complexity, which relaxes the notion of tractability as polynomial-time solvability to allow non-polynomial dependence on certain problem-specific parameters. A whole new set of interesting questions arises if we look at CSPs from this point of view. Most CSP dichotomy questions can be revisited by defining a parameterized version; so far, very little work was done in this direction compared to investigations in classical complexity. A new research direction (often called "parameterizing above the guaranteed bound") led to unexpected positive results for Max r-SAT by Alon *et al.* in 2010. In this direction, the basic question is to decide the fixed-parameter tractability of the following type of problems: if some easily computable estimate guarantees satisfaction at least E constraints, find an assignment that satisfies at least $E+k$ constraints. Y. Makarychev presented recent results, including approximation issues, in this direction that concern the so-called ordering CSP. Wahlström and Yoshida described how algorithms for this problem, also for VCSP, can be designed when the estimate is given by the Linear Programming relaxation.

Logic and the complexity of CSP. Starting from earlier work by Kolaitis ad Vardi, concepts and techniques from logic have provided unifying explanations for many tractable CSPs. This has led to the pursuit of classifications of CSP with respect to *descriptive complexity*, i.e. definability in a given logic. Logics considered in this context include first order logic and its extensions, finite-variable logics, the logic programming language Datalog and its fragments. Kazda presented his recent results on the two most important open problems on descriptive complexity of CSPs, where he showed that one of these problems reduces to the other. These results are also related to dichotomy questions for complexity classes L and NL.

The CSP can be recast as the problem of deciding satisfiability of existential conjunctive formulas. Chen described recent results in this direction that also involve counting and parameterised complexity. Natural extension of this framework that also allows universal quantifiers is known as the Quantified CSP (QCSP). New results on the complexity of language-restricted QCSPs were presented by Martin. Zhuk gave a proof of an algebraic result that has direct strong consequences for complexity classification of QCSPs.

Bodirsky and Pinsker presented latest developments in infinite domain CSPs, obtained via a mixture of model-theoretic and algebraic methods.

Ochremiak investigated finite-domain CSPs on infinite instances definable by formulas in first-order logic.

Fig. 7.16
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.42 Digital Scholarship and Open Science in Psychology and the Behavioral Sciences

Organizers: Alexander Garcia Castro, Janna Hastings, Robert Stevens, and Erich Weichselgartner

Seminar No. 15302

Date: July 19–24, 2015 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.5.7.42

Participants: Xavier Aimé, Dietrich Albert, George Alter, Björn Brembs, Mike Conlon, Oscar Corcho, Susann Fiedler, Alexander Garcia Castro, Paul Groth, William Gunn, Janna Hastings, Caroline Jay, Iris-Tatjana Kolassa, Silvia Koller, Christoph Lange, Maryann Martone, Russell Poldrack, Alec Smecher, Daniel Staemmler, Robert Stevens, Gary VandenBos, Hal Warren, Erich Weichselgartner

Researchers across many domains have invested significant resources to improve transparency, reproducibility, discoverability and, in general, the ability to share and empower the community. Digital Scholarship and Open Science are umbrella terms for the movement to make scientific research, its tools and data and dissemination accessible to all members of an inquiring society, amateur or professional. Digital infrastructures are an essential prerequisite for such open science and digital scholarship; the biomedical domain illustrates this culture. An impressive digital infrastructure has been built; this allows us to correlate information from genomes to diseases, and, by doing so, to support movements such as panomic studies and personalized medicine. A high degree of interdisciplinary work was necessary in building this infrastructure; the large quantities of data being produced, the high degree of interrelatedness, and, most of all, the need for this mingling of many types of data in a variety of forms forged this collaboration across the community and beyond.

The Behavioral Sciences, comprising psychology but also psychobiology, criminology, cognitive science and neuroscience, are also producing data at significant rates; by the same token, understanding mental health disorders requires correlating information from diverse sources – e.g. cross-referencing clinical, psychological, and genotypic sources. For example, flagship projects such as the Brain Activity Map (BAM, also known as the BRAIN initiative[46]) are generating massive amounts of data with potential benefit to mental health and psychology; conversely, projects like BAM could benefit from information currently being generated by psychologists. Our ability to make continued progress in understanding the mind and brain depends on finding new ways to organize and synthesize an ever-expanding body of information.

The *'Digital Scholarship and Open Science in Psychology and the Behavioral Sciences'* Dagstuhl Perspectives Workshop was conceived with one problem in mind: that of facilitating the construction of an integrative infrastructure in Psychology and Behavioral Sciences. The motivation for this workshop was to *'foster the discussion around the problem of understanding the Web as an integrative platform, and how e-science can help us to do better research.'* With these points in mind, we gathered an interdisciplinary group of experts, including computer scientists, psychologists and behavioral scientists. In their research, they are addressing issues in data standards, e-science, ontologies, knowledge management, text mining, scholarly communication, semantic web, cognitive sciences, neurosciences, and psychology. Throughout the Workshop, this group worked on devising a roadmap for building such an interoperability layer.

The seminar started with a number of keynote sessions from well-known authorities in each area to introduce the necessary background and form a common baseline for later discussions. A core theme that emerged was the cross-domain challenge in establishing a common language. We jointly undertook the effort to define an integrative scenario illustrating how digital infrastructures could help psychologists and behavioral scientists to do research that takes advantage of the new digital research landscape. In order to achieve this, the computational scientists needed to better understand the current working practices of the psychologists. For instance, the nature and structure of their data and experiments; moreover, computer scientists needed to understand the flow of information, from the conception of an idea, through defining a study plan, executing it and finally having the investigation published. They learned that the work of psychologists and behavioral scientists strongly relies on questionnaires and experiments as ways of collecting data, and on statistics as a tool for analyzing data, and that the replicability of experiments

[46] http://www.braininitiative.nih.gov/

is a key concern. In a similar vein, psychologists and behavioral scientists needed concrete examples illustrating how computer science enables FAIR (= findable, accessible, interoperable and reusable) infrastructures that allow researchers to discover and share knowledge – bearing in mind data protection issues.

Two break-out groups were organized. The purpose was to have a full picture of digital scholarship in action when applied to psychology and behavioral investigations, most importantly e-science assisting researchers in sharing, discovering, planning and running investigations. The full research life cycle had to be considered. Both groups worked up their respective scenarios independently. The visions were then exchanged in an inter-group meeting. Interestingly, various issues arose when discussing the specifics from each vision for digital scholarship; for instance, the importance of understanding scholarly communication beyond the simple act of getting one's results published. Furthermore, the need to integrate tools into platforms where researchers could openly register their projects and plan and manage their workflows, data, code and research objects, was extensively discussed. Within this framework, the need for controlled vocabularies, standards for publishing and documenting data and metadata, persistent identifiers for datasets, research objects, documents, organizations, concepts and people, open APIs to existing services and instruments, and reporting structures were understood; these elements were articulated in the examples where the researchers and research were at the center of the system. Discussions also addressed fears in the community and thus the need to open up the current research landscape in small steps.

The seminar proved to be a fertile discussion field for interdisciplinary collaborations and research projects across previously disparate fields with the potential of significant impact in both areas. The need for a digital infrastructure in psychology and behavioral sciences was accepted by all the attendants; communicating this message with a clear implementation vision to funding agencies, professional societies and the community in general was identified as a key priority. It was decided that we needed another meeting in 2016; during that follow-up, the emphasis should be on developing a research agenda. As this is a relatively new topic in psychology and behavioral sciences, it was also decided to contact publishers and professional organizations, e.g. the Sloan Foundation, the APA and the APS, and work with them in conveying the message about increasing openness. If we want to understand how cognition is related to the genome, proteome and the dynamics of the brain, then interoperability, data standards and digital scholarship have to become a common purpose for this community. Funding has to be made available, initially for an assessment of the uptake of existing key resources and infrastructures, and then for implementing further Digital Scholarship and Open Science infrastructures as well as for building the skills in a community that is not yet widely familiar with the relevant enabling technologies. Finally, once sufficient technical support is in place, sustainable incentives for sharing research objects should be put in practice.

7.43 Power-Bounded HPC Performance Optimization

Organizers: Dieter Kranzlmüller and Barry L. Rountree

Seminar No. 15342

Date: August 16–21, 2015 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.5.8.1

Participants: Michael Bader, Natalie Bates, Pete Beckman, Jonathan Eastep, Neha Gholkar, Joseph Greathouse, Thomas Ilsche, Dieter Kranzlmüller, Stephanie Labasan, David K. Lowenthal, Matthias Maiterth, Andres Marquez, Yousri Mhedheb, Shirley V. Moore, Frank Mueller, Andreas Raabe, David Richards, Suzanne Rivoire, Barry L. Rountree, Martin Schulz, Kathleen Sumiko Shoga, Torsten Wilde

The Dagstuhl Perspectives Workshop 15342 "Power-Bounded HPC Performance Optimization" has been an interesting experience, as in contrast to other workshops, we focused on the unknown characteristics of future exascale systems rather than on the state-of-the-art of todays petascale architectures. In order to do this, a large fraction of the workshop was spent on in-depth discussions in three working groups, while plenary sessions served to provide impulses on specific topics and to synthesize the findings of the breakout sessions. The key ingredient of this workshop has been the interaction between the participants, leading to several new collaborations across vendors, national laboratories and academia.

The key findings of the workshop can be identified as follows:

- Power-bound performance optimization has different objectives according to the respective targets and operational goals. While infrastructure providers are often bound to a specific spending, users want to utilize a resource at the maximum of its capabilities. As a result, any power-bound optimization must address multiple criteria, and the solution is rarely straight-forward but specific for a given setting.
- The currently available information on each layer of the computing environment is insufficient. Both, the availability of information with respect to its power characteristics, as well as the exchange between different layers, needs to be improved in order to optimize the operation of infrastructures and the execution of applications on a given system.
- Due to the number of dependencies, any optimization needs to find a good balance between "user happiness", total costs, and performance. These characteristics are important for both, providers and users, and a careful balancing strategy needs to be implemented without harming any interests of the actors too much.

The discussions at the Dagstuhl Perspectives Workshop have lead to the identification of a number of technical problems, which need to be addressed in the near future before achieving optimal results in a power-bound environment. As a conclusion, the participants agreed that a strategic and tactical agenda is needed, which identifies the individual problems and technologies as well as their interconnections, such that future systems can utilize this knowledge for new approaches of power-bound HPC performance optimization. The results of this investigations should be made available as a white book, which describes the strategy for future exascale systems.

7.44 Computational Mass Spectrometry

Organizers: Rudolf Aebersold, Oliver Kohlbacher, and Olga Vitek

Seminar No. 15351

Date: August 23–28, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.8.9

Participants: Rudolf Aebersold, Theodore Alexandrov, Naomi Altman, Nuno Bandeira, Andreas Beyer, Sebastian Böcker, Robert Chalkley, Joshua Elias, Laurent Gatto, Anne-Claude Gingras, Matthias Gstaiger, Kasper Daniel Hansen, Henning Hermjakob, Michael Hoopmann, Lukas Käll, Oliver Kohlbacher, Bernhard Küster, Kathryn Lilley, Lennart Martens, Sven Nahnsen, Pedro José Navarro Alvarez, Robert Ness, Jonathon O'Brien, Patrick Pedrioli, Knut Reinert, Lukas Reiter, Bernhard Renard, Hannes Röst, Karen Sachs, Timo Sachsenberg, Albert Sickmann, Stephen Tate, Stefan Tenzer, Michael L. Tress, Olga Vitek, Christine Vogel, Susan T. Weintraub, David Wishart, Bernd Wollscheid, Nicola Zamboni

Motivation

Mass Spectrometry (MS) is an extremely flexible analytical technique, with applications ranging from crime lab investigations to testing to disease biomarkers in a clinic. The publication of the first human genome in 2001 was a key event that lead to the application of mass spectrometry to map out the human proteome, and later the human metabolome; i.e. all the biomolecules encoded in the genome that constitute biological function. The result was the creation of a tremendous amount of spectrometric data and a dearth of tools for data analysis, motivating the development of computational tools. The tool developers came from several expert domains; life scientists applying mass spectrometry built tools to automate their new workflows, analytical chemists and engineers developing the instruments built software to analyze devise measurements; network and database infrastructure professionals built resources for storing and sharing data in the cloud, and bioinformaticians and statisticians developed algorithms and statistical methods for data analysis. There is an ongoing need for the different disciplines to learn each other's languages, make tools interoperable, and establish common goals for development.

Goals

The seminar 'Computational Mass Spectrometry' is a follow-up seminar to the successful Dagstuhl seminars on 'Computational Proteomics' and 'Computational Mass Spectrometry' (05471, 08101 and 14371).

The seminar aimed at bringing together scientists from a wide range of backgrounds and identify open issues and future research directions in computational mass spectrometry.

Results

Already on the first days the seminar resulted in very lively discussions. The time allotted to the introductory talks had to be expanded to account for this. The discussions sparked off during the introductory talks led to the formation of several working groups. These groups formed and re-formed on demand, also based on discussion on the previous evenings. The full report documents the discussions and results in these groups through the notes taken. Some of these discussion (e.g., the one on false discovery rates) was of interest to all participants and took place as plenary discussions in the large lecture hall. Other discussions were more focussed and thus had a smaller number of participants.

Some of the discussion will certainly lead to joint research participants. A first tangible outcome is a joint paper already accepted in the *Journal of Proteome Research* (L. Gatto, K. D. Hansen, M. R. Hoopmann, H. Hermjakob, O. Kohlbacher, A. Beyer, "*Testing and validation of computational methods for mass spectrometry*," DOI: 10.1021/acs.jproteome.5b00852) on benchmarking and validating computational methods for mass spectrometry. This working group developed conceptual ideas for benchmarking algorithms and implemented a web-based repository holding (http://compms.org/RefData) benchmark datasets that will hopefully make comparison of algorithms more transparent in the future. We are confident that the discussions of other working groups and the contacts made during the evening hours in Dagstuhl will result in many more collaborations and publications in the future.

The field of computational mass spectrometry is rapidly evolving. Participants identified a wide range of challenges arising from technological developments already at the horizon but also from the broadening on the application side. We thus intend to revisit the field in the coming years in a Dagstuhl seminar again, most likely organized by different leaders of the field in order to account for these upcoming changes.

7.45 Design of Microfluidic Biochips: Connecting Algorithms and Foundations of Chip Design to Biochemistry and the Life Sciences

Organizers: Krishnendu Chakrabarty, Tsung-Yi Ho, and Robert Wille

Seminar No. 15352

Date: August 23–26, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.8.34

Participants: Mirela Alistar, Krishnendu Chakrabarty, Rolf Drechsler, William H. Grover, Tsung-Yi Ho, Juinn-Dar Huang, Oliver Keszöcze, Bing Li, Pietro Lio', Jan Madsen, Sebastian Maerkl, Paul Pop, Sudip Roy, Ulf Schlichtmann, Kwanwoo Shin, Rüdiger Trojok, Steve Wereley, Robert Wille, Hailong Yao

Advances in microfluidic technologies have led to the emergence of biochip devices for automating laboratory procedures in biochemistry and molecular biology. These devices enable the precise control of nanoliter-scale biochemical samples and reagents. Therefore, *Integrated Circuit* (IC) technology can be used to transport a "chemical payload" in the form of micro- or nano-fluidic carriers such as droplets. As a result, non-traditional biomedical applications and markets (e.g., high-throughput DNA sequencing, portable and point-of-care clinical diagnostics, protein crystallization for drug discovery), and fundamentally new uses are opening up for ICs and systems. This represents a More than Moore-approach.

Miniaturized and low-cost biochip systems are revolutionizing a diverse range of applications, e.g., air quality studies, point-of-care clinical diagnostics, drug discovery, and DNA sequencing. Frost & Sullivan recently predicted a 13.5% Compound Annual Growth Rate for the US biochip ("lab-on-chip") market during 2008-2015, and the market size for lab-on-chip alone (not including microarrays, biosensors, and microreactors) is expected to be over $1.6 billion in 2015. Similar growth is anticipated in other parts of the world, especially Europe and Japan. On a broader scale, the annual US market alone for in vitro diagnostics is as high as $10 billion and similar figures have been estimated for the drug discovery market. For clinical diagnostics, it has been predicted that we will soon see 15 billion diagnostic tests/year worldwide.

However, continued growth (and larger revenues resulting from technology adoption by pharmaceutical and healthcare companies) depends on advances in chip integration and design-automation tools. Thus, there is a need to deliver the same level of *Computer-Aided Design* (CAD) support to the biochip designer that the semiconductor industry now takes for granted. In particular, these CAD tools will adopt computational intelligence for the optimization of biochip designs. Also, the design of efficient CAD algorithms for implementing biochemistry protocols to ensure that biochips are as versatile as the macro-labs that they are intended to replace. This is therefore an opportune time for the software and semiconductor industry and circuit/system designers to make an impact in this emerging field.

Recent years have therefore seen growing interest in design methods and design-automation tools for the digital microfluidic platform, with special issues of *IEEE Transactions on CAD and IEEE Design & Test of Computers*, special sessions at *DAC*, *ISPD*, *ASPDAC*, and *ICCAD*, and workshops/tutorials at *ISCAS*, *ICCAD*, *SOCC*, and *DATE*. A number of CAD research groups worldwide (e.g., Duke University; Carnegie Mellon University; University of Texas at Austin; Rensselaer Polytechnic University; University of California at Riverside; University of Washington; Technical University of Denmark; Technische Universität München; University of Bremen; National Tsing Hua University; National Chiao Tung University, National Taiwan University; Tsinghua University; Indian Statistical Institute; Ritsumeikan University; Nanyang Technological University; Johannes Kepler University Linz) have initiated research projects on CAD for microfluidic biochips.

The goal of the seminar was to bring together experts in order to present and to develop new ideas and concepts for the design automation algorithms and tools for microfluidic biochips. Areas ranging from architecture, synthesis, optimization, verification, testing, and beyond have been covered. Topics which have been discussed included besides others:

- Architectural synthesis
- Behavior-level synthesis
- Cooling for integrated circuits
- Cross-contamination removal
- Cyberphysical integration
- Device modeling
- Drug-delivery biochips

- Fault modeling, testing, and protocol verification
- Light-actuated biochips
- Numerical simulation
- On-chip sensors
- Paper-based microfluidics
- Particle microfluidics
- Physical design
- Pin-constrained design
- Sample preparation

As results we received a better understanding of the respective areas, new impulses for further research directions, and ideas for areas that will heavily influence research in the domain of design automation on microfluidic biochips within the next years. The seminar facilitated greater interdisciplinary interactions between chip designers, bioengineers, biochemists, and theoretical computer scientists.

The high-quality presentations and lively discussions have been ensured by carefully selected experts who participated at the seminar. All of them have established for themselves a stellar reputation in the respective domains. While researchers working on design automation and optimization of microfluidic biochips build the majority of the participants, also some experts from surrounding research areas attended. For example, researchers working on emerging architectures and applications of microfluidic biochips provided the needed insight for the discussions about the practical problem formulation for commercialized product. Computer scientists with a focus on computer-aided design enriched the discussions about the top-down design methodology and optimization of large-scale components like mixers and routing channels. Therewith, the unique concept of Dagstuhl seminars was applied in order to bring researchers from different domains together so that the interdisciplinary topics could have been discussed and progress in these areas has been made.

7.46 Mathematical and Computational Foundations of Learning Theory

Organizers: Matthias Hein, Gabor Lugosi, and Lorenzo Rosasco

Seminar No. 15361

Date: August 30 to September 4, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.8.54

Participants: Shivani Agarwal, Animashree Anandkumar, Peter L. Bartlett, Shai Ben-David, Gilles Blanchard, Stephane Boucheron, Sebastien Bubeck, Joachim M. Buhmann, Constantine Caramanis, Sou-Cheng Choi, Luc Devroye, Jack Fitzsimons, Antoine Gautier, Remi Gribonval, László Györfi, Moritz Hardt, Matthias Hein, Prateek Jain, Stefanie Jegelka, Felix Krahmer, Andreas Krause, Lek-Heng Lim, Gabor Lugosi, Robert D. Nowak, Guillaume Obozinski, Duy Khanh Pham, Lorenzo Rosasco, Alessandro Rudi, Sivan Sabato, Karin Schnass, Dejan Slepcev, Nathan Srebro, Yannik Stein, Alexandre Tsybakov, Ruth Urner, Silvia Villa, Rachel Ward, Colin White, Robert C. Williamson, Yuan Yao, Ding-Xuan Zhou

Machine learning is nowadays a central field in computer science. Over the last decade the statistical learning approach has been successfully applied in many areas such as bioinformatics, computer vision, robotics and information retrieval. We believe that the main reasons for the success of machine learning are its strong theoretical foundations and its multidisciplinary approach integrating aspects of computer science, applied mathematics, and statistics among others.

Two very successful conferences titled "Mathematical Foundations of Learning Theory" in Barcelona 2004 and Paris 2006 have been inspired by this point of view on the foundations of machine learning. In 2011 the Dagstuhl seminar "Mathematical and Computational Foundations of Learning Theory" has been organized in the same spirit, bringing together leading researchers from computer science and mathematics to discuss the state of the art and future challenges in machine learning. The 2011 Dagstuhl seminar has been the first to cover a wide range of facets of modern learning theory and has been unanimously considered a success by the participants. Since 2011 new challenges have emerged largely motivated by the availability of data-sets of unprecedented size and complexity. It is now common in many applied domains of science and technology to have datasets with thousands and even millions data-points, features and attributes/categories. For example ImageNet (http://image-net.org) is a computer vision database for object recognition including one million images of one thousands different objects, and image representations are often of the order of several tens of thousands features. Datasets of analogous complexity are customary in biology and information science (e.g. text classification). The need of analyzing and extracting information from this kind of data has posed a host of new challenges and open questions.

The second Dagstuhl seminar on "Mathematical and Computational Foundations of Learning Theory" covered broadly recent developments in the area of learning. The main focus was on two topics:

- **Interplay between Optimization and Learning**
  While statistical modeling and computational aspects have for a long time been considered separate steps in the design of learning algorithms, dealing effectively with big data requires developing new strategies where statistical and computational complexities are taken simultaneously into account. In other words, the trade-off between optimization error and generalization error has to be exploited. On the other hand it has very recently been noticed that several non-convex NP-hard learning problems (sparse recovery, compressed sensing, dictionary learning, matrix factorization etc.) can be solved efficiently and optimally (in a global sense) under conditions on the data resp. the chosen model or under the use of additional constraints.
- **Learning Data Representations**
  Data representation (e.g. the choice of kernels or features) is widely acknowledged to be the crucial step in solving learning problems. Provided with a suitable data representation, and enough labeled data, supervised algorithms, such as Support Vector Machines or Boosting, can provide good generalization performance. While data representations are often designed ad hoc for specific problems, availability of large/huge amount of unlabeled data have recently motivated the development of data driven techniques, e.g. dictionary learning, to adaptively solve the problem. Indeed, although novel tools for efficient data labeling have been developed (e.g. Amazon Mechanical Turk– http://mturk.com) most available data are unlabeled and reducing the amount of (human) supervision needed to effectively solve a task remains an important open challenge. While up-to-now the theory of supervised learning has become a mature field, an analogous theory of unsupervised and semi-supervised

learning of data representation is still in its infancy and progress in the field is often assessed on a purely empirical basis.

The seminar featured a series of talks on both topics with interesting and exciting new results which lead to insights in both areas as well as a lot of discussion and interaction between the participants which for sure will manifest in several follow-up papers. Also it became obvious during the seminar that there are close connections between these two topics. Apart from these two main topics several other aspects of learning theory were discussed, leading to a quite complete picture on the current state-of-the-art in the field.

Acknowledgements. We would like to thank Dagmar Glaser and the staff at Schloss Dagstuhl for their continuous support and great hospitality which was the basis for the success of this seminar.

7.47 Present and Future of Formal Argumentation

Organizers: Dov M. Gabbay, Massimiliano Giacomin, Beishui Liao, and Leendert van der Torre

Seminar No. 15362

Date: August 30 to September 4, 2015 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.5.8.74

Participants: Pietro Baroni, Ringo Baumann, Stefano Bistarelli, Alexander Bochman, Gerhard Brewka, Katarzyna Budzynska, Martin Caminada, Federico Cerutti, Wolfgang Dvorak, Dov M. Gabbay, Massimiliano Giacomin, Tom Gordon, Beishui Liao, Henry Prakken, Chris Reed, Odinaldo Rodrigues, Guillermo R. Simari, Matthias Thimm, Leendert van der Torre, Bart Verheij, Emil Weydert, Stefan Woltran

Diverse kinds of reasoning and dialogue activities can be captured by argumentation models in a formal and still quite intuitive way, thus enabling the integration of different specific techniques and the development of applications humans can trust. Formal argumentation lays on solid bases, such as extensively studied theoretical models at different levels of abstraction, efficient implementations of these models, as well as a variety of experimental studies in several application fields. In order to be able to convert the opportunities of the present into actual results in the future, the formal argumentation research community needs however to reflect about the current assets and weaknesses of the field and to identify suitable strategies to leverage the former and to tackle the latter. As an example, the definition of standard modeling languages and of reference sets of benchmark problems are still in their infancy, reference texts for newcomers are missing, the study of methodological guidelines for the use of theoretical models in actual applications is a largely open research issue.

The goal of this Dagstuhl Perspectives Workshop was to gather the world leading experts in formal argumentation in order to develop a SWOT (Strength, Weaknesses, Opportunities, Threats) analysis of the current state of the research in this field and to draw accordingly some strategic lines to ensure its successful development in the future.

The Perspectives Workshop was held between August 30 to September 4, 2015, with 22 participants from 10 countries. With the aim of developing a critical survey of the field for the argumentation community and for potential newcomers, the organizers agreed to assemble a handbook of formal argumentation, and encouraged participants to present their view on different topics in the area. Besides individual presentations, the program included collective discussions on general issues arising from individual presentations, as well as working groups.

Individual presentations concerned introductory overviews, logical problems and requirements for formal argumentation, specific formalisms and methodologies, relationship between different approaches and applications. While a limit of half an hour per talk was initially established, we decided to leave the time for discussion relatively open, since several open topics and new developments were envisaged out of presentations.

Collective discussions have been arranged along four topics, i.e. basic concepts and foundations, specific formalisms for argumentation, algorithms, and connections both inside the argumentation field and with outside research topics.

We organized three discussion groups each headed by one organizer. Each group was asked to identify the most important open problems in argumentation. Interestly enough, there was little intersection between the three outcomes, i.e. the three groups came out with different problems. Many of them concerned foundational issues of the theory, e.g, how to formally represent various kinds of arguments and how to identify sets of postulates on the reasoning activity over arguments in specific contexts. On the other hand, the relationship between argumentation and other research fields (e.g. natural language processing, machine learning, human computer interaction, social choice) was seen to be of major importance, especially to develop more applications.

The unique setting and atmosphere of Dagstuhl provided the ideal environment to exchange ideas on future directions of argumentation, with discussions often lasting all the evening and the first part of the night.

The Perspectives Workshop concluded with the presentation of the results yielded by the group discussions, that in our opinion will lead to collaborative research, scientific papers and funded international projects in the future.

Fig. 7.17
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of TreeState Productions GmbH, www.treestate.de.

7.48 Quantum Cryptanalysis

Organizers: Michele Mosca, Martin Roetteler, Nicolas Sendrier, and Rainer Steinwandt

Seminar No. 15371

Date: September 6–11, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.1

Participants: Gorjan Alagic, Aleksandrs Belovs, Daniel J. Bernstein, Jean-François Biasse, Alexei Bocharov, Harry Buhrman, André Chailloux, Jintai Ding, Hang Dinh, Jürgen Eschner, Jennifer Katherine Fernick, Tommaso Gagliardoni, Markus Grassl, Sean Hallgren, Peter Hoyer, Andreas Hülsing, Stacey Jeffery, Stavros Kousidis, Thijs Laarhoven, Bradley Lackey, Tanja Lange, Anthony Leverrier, Yi-Kai Liu, Alexander May, Kirill Morozov, Michele Mosca, Michael Naehrig, Maris Ozols, Ray Perlner, Martin Roetteler, Christian Schaffner, John M. Schanck, Claus-Peter Schnorr, Nicolas Sendrier, Dan J. Shepherd, Daniel Smith-Tone, Fang Song, Rainer Steinwandt, Krysta Svore, Tsuyoshi Takagi, Enrico Thomae, Jean-Pierre Tillich, Joop van de Pol, Frank K. Wilhelm, Bo-Yin Yang

It is known that quantum algorithms exist that jeopardize the security of most of our widely-deployed cryptosystems, including RSA and Elliptic Curve Cryptography. It is also known that advances in quantum hardware implementations are making it increasingly likely that large-scale quantum computers will be built in the near future that can implement these algorithms and devastate most of the world's cryptographic infrastructure. What is not known is an estimate of the resources that will be required to carry out these attacks – or even whether other quantum attacks exist that have not yet been accounted for in our security estimates. In this seminar, we examined both computational resource estimates for meaningful quantum cryptanalytic attacks against classical (i.e.: conventional) cryptography, as well as the security of proposed quantum-safe cryptosystems against emerging quantum cryptanalytic attacks.

This seminar had a number of research highlights spanning the areas of implementations of quantum hardware and software, quantum algorithms, and post-quantum cryptography.

Implementations of quantum information processing were outlined to help contextualize the current state of quantum computation. Recent advances in the synthesis of efficient quantum circuits were presented, as well as an update on implementations – particularly in the domain of superconducting integrated circuits. Seminar participants were warned that traditional approaches to the modeling of quantum processors may be reaching an end, while the LIQUi|> software architecture for control of quantum hardware and simulation of quantum algorithms was unveiled. Challenges involving practical costs for error correction in systems with specific types of quantum memory (particularly quantum bucket brigade RAM architectures) were articulated.

In the domain of algorithmic advances, seminar participants demonstrated quantum improvements on the gapped group testing problem, as well as improvements on lattice sieving using nearest neighbour search algorithms. A discussion of how quantum computers can sometimes provide quadratic speedup for the differential cryptanalysis of symmetric-key cryptosystems was also presented. A quantum version of the unique-SVP algorithm was discussed, but it was found to have slightly worse performance than its' classical counterpart. For the purposes of improving our understanding quantum algorithms before large-scale quantum computers become available, a technique involving trapdoor simulation of quantum algorithms was proposed.

Seminar participants also gave a number of recent results in the domain of quantum-safe cryptography. These included a provably-secure form of Authenticated Key Exchange based on the Ring-Learning with Errors problem, a proposal for a quantum-safe method to prevent key leakage during key agreement failure stemming from invalid public keys, and updates on hash-based digital signatures. The EU PQCRYPTO project also presented some preliminary recommendations for post-quantum cryptography.

In the domain of code-based cryptography, it was demonstrated that assuming hardness of Niederreiter problem, CFS signatures are strongly existentially unforgeable in the random oracle model. A number of results related to lattice reduction were also presented, including an improvement on the BKZ lattice reduction algorithm, some lattice enumeration work involving factoring integers by CVP algorithms for the prime number lattice, and a reduction of gapped uSVP to the Hidden Subgroup Problem in dihedral groups. A LIQUi|> implementation of a quantum algorithm to extract hidden shift was also presented, as well as demonstration of instances of HSPs over dihedral group which can be efficiently solved on a quantum computer. Seminar participants also proposed alternative ways of thinking about the dihedral coset problem, including some hardness reductions. A very new result on finding a generator of a principal ideal was also debuted at this seminar and provoked lively and ongoing discussion among participants.

Other talks were presented on diverse and compelling topics including quantum-mechanical means for program obfuscation, and a means for quantum indistinguishability of some types of ciphertext messages. A presentation was also made about how standardization bodies and industry deal with information security and risk, and many discussions – both formal and informal – among participants began to deal with the applied challenges of developing and deploying quantum-safe information security standards and tools.

Overall, the success of this seminar can be observed not only through the quantity of new results, but also in their diversity and interdisciplinarity. While there exist venues for cryptography and cryptanalysis, for quantum algorithms, and for implementations of quantum information processing, it remains critical that these communities continue to come together to ensure rigorous and broad cryptanalysis of proposed quantum-safe cryptographic algorithms, and to share a well-defined mutual understanding of the quantum-computational resource requirements – and their present availability – for attacking both public and symmetric key cryptography. The security of the world's information depends on it.

The organizers (Michele Mosca, Martin Roetteler, Nicolas Sendrier, and Rainer Steinwandt) are grateful to the participants of this seminar and the team of Schloss Dagstuhl for an inspiring and productive third edition of this seminar series.

7.49 Information from Deduction: Models and Proofs

Organizers: Nikolaj S. Bjørner, Jasmin Christian Blanchette, Viorica Sofronie-Stokkermans, and Christoph Weidenbach

Seminar No. 15381

Date: September 13–18, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.18

Participants: Noran Azmy, Franz Baader, Peter Baumgartner, Christoph Benzmüller, Nikolaj S. Bjørner, Jasmin Christian Blanchette, Martin Bromberger, Catherine Dubois, Bruno Dutertre, Carsten Fuhs, Silvio Ghilardi, Jürgen Giesl, Alberto Griggio, Arie Gurfinkel, Liana Hadarean, Reiner Hähnle, Matthias Horbach, Swen Jacobs, Dejan Jovanovic, Deepak Kapur, George Karpenkov, Zachary Kincaid, Konstantin Korovin, Christopher Lynch, Aart Middeldorp, Tobias Nipkow, Albert Oliveras, Andrei Paskevich, Alexander Rabinovich, Giles Reger, Albert Rubio, Andrey Rybalchenko, Stephan Schulz, Viorica Sofronie-Stokkermans, Christian Sternagel, Geoff Sutcliffe, Cesare Tinelli, Andrei Voronkov, Christoph Weidenbach, Sarah Winkler, Burkhart Wolff, Jian Zhang

Models and proofs are the quintessence of logical analysis and argumentation. Many applications of deduction tools need more than a simple answer whether a conjecture holds; often additional information – for instance proofs or models – can be extremely useful. For example, proofs are used by high-integrity systems as part of certifying results obtained from automated deduction tools, and models are used by program analysis tools to represent bug traces. Most modern deductive tools may be trusted to also produce a proof or a model when answering whether a conjecture is a theorem or whether a certain problem formalized in logic has a solution. Moreover, major progress has been obtained recently by procedures that rely on refining a simultaneous search for a model and a proof. Thus, proofs and models help producing models and proofs, and applications use proofs and models in many crucial ways.

Below, we point out several directions of work related to models and proofs in which there are challenging open questions:

- *Extracting proofs from derivations.* An important use of proof objects from derivations is for applications that require certification. But although the format for proof objects and algorithms for producing and checking them has received widespread attention in the research community, the current situation is not satisfactory from a consumer's point of view.
- *Extracting models from derivations.* Many applications rely on models, and models are as important to certify non-derivability. Extracting models from first-order saturation calculi is a challenging problem: the well-known completeness proofs of superposition calculi produce perfect models from a saturated set of clauses. The method is highly non-constructive, so extracting useful information, such as "whether a given predicate evaluates to true or false under the given saturated clauses," is challenging. The question of representation is not yet well addressed for infinite models.
- *Using models to guide the search for proofs and vice versa.* An upcoming next generation of reasoning procedures employ (partial) models/proofs for proof search. They range from SAT to first-order to arithmetic reasoning and combinations thereof. It remains an open question what properties of models are crucial for successful proof search, how the models should be dynamically adapted to the actual problem, and how the interplay between the models and proof search progress through deduction should be designed.
- *External applications of models and proofs.* Models and proofs are used in various ways in applications. So far application logics and automated proof search logics have been developed widely independently. In order to get more of a coupling, efforts of bringing logics closer together or the search for adequate translations are needed.

This Dagstuhl seminar allowed to bring together experts for these topics and invited discussion about the production and consumption of proofs and models. The research questions pursued and answered include:

- To what extent is it possible to design common exchange formats for theories, proofs, and models, despite the diversity of provers, calculi, and formalisms?
- How can we generate, process, and check proofs and models efficiently?
- How can we search for, represent, and certify infinite models?
- How can we use models to guide proof search and proofs to guide model finding?
- How can we make proofs and models more intelligible, yet at the same time provide the level of detail required by certification processes?

7.50 Modeling and Simulation of Sport Games, Sport Movements, and Adaptations to Training

Organizers: Ricardo Duarte, Björn Eskofier, Martin Rumpf, and Josef Wiemeyer

Seminar No. 15382

Date: September 13–16, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.38

Participants: Arnold Baca, Eva Dorschky, Ricardo Duarte, Stefan Endler, Björn Eskofier, Irfan A. Essa, Hugo Folgado, Katrin Hoffmann, Anne Danielle Koelewijn, John Komar, Martin Lames, Roland Leser, Daniel Link, Jim Little, Stuart Morgan, Bernhard Moser, Jürgen Perl, Robert Rein, Karen Roemer, Martin Rumpf, Tiago Guedes Russomanno, Dietmar Saupe, Heiko Schlarb, Malte Siegle, Michael Stöckl, Antonie van den Bogert, Anna Volossovitch, Hendrik Weber, Josef Wiemeyer

Computational modeling and simulation are essential to analyze human motion and interaction in sport science, sport practice and sport industry. Applications range from game analysis, issues in exercising like training load-adaptation relationship, motor control and learning, to biomechanical analysis. New challenges appear due to the rapid development of information and communication technologies (ICT) as well as the enormous amount of data being captured within training and competition domains. The motivation of this seminar was to enable an interdisciplinary exchange between sports and computer scientists as well as sport practice and industry to advance modeling and simulation technologies in selected fields of applications: sport games, sport movements and adaptations to training.

From September 13 to September 16, 2015 about 29 representatives of science, practice and industry met at the Leibniz-Zentrum für Informatik in Schloss Dagstuhl to discuss selected issues of modelling and simulation in the application fields of sport games, sport movements and adaptation to training. This seminar was the fifth in a series of seminars addressing computer science in sport, starting in 2006. Based on previously selected issues, four main streams were identified:

- Validation and model selection
- Sensing and tracking
- Subject-specific modelling
- Training and sport games

The talks addressing these four topics are summarized in the full report. They have been arranged according to the three main application fields: sport games, sport movements, and adaptations to training. In addition, generic comments on modeling in industry and science are presented. Moreover, the final discussion is summarized and a conclusion of the seminar is drawn in the full report.

7.51 Algorithms and Complexity for Continuous Problems

Organizers: Aicke Hinrichs, Joseph F. Traub, Henryk Woźniakowski, and Larisa Yaroslavtseva

Seminar No. 15391

Date: September 20–25, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.57

Participants: James M. Calvin, Ronald Cools, Sonja Cox, Steffen Dereich, Stefan Geiss, Michael Gnewuch, Mario Hefter, Stefan Heinrich, Aicke Hinrichs, Alexander Keller, Peter Kritzer, Thomas Kühn, Peter Mathé, Thomas Müller-Gronbach, Andreas Neuenkirch, Dong Nguyen, Erich Novak, Dirk Nuyens, Jens Oettershagen, Sergei Pereverzyev, Leszek Plaskota, Pawel Przybylowicz, Holger Rauhut, Klaus Ritter, Daniel Rudolf, Winfried Sickel, Pawel Siedlecki, Ian H. Sloan, Jeremy Staum, Ingo Steinwart, Mario Ullrich, Tino Ullrich, Markus Weimar, Larisa Yaroslavtseva, Marguerite Zani

This was already the 12th Dagstuhl Seminar on Algorithms and Complexity for Continuous Problems over a period of 24 years. It brought together researchers from different communities working on computational aspects of continuous problems, including computer scientists, numerical analysts, applied and pure mathematicians. Although the seminar title has remained the same, many of the topics and participants change with each seminar and each seminar in this series is of a very interdisciplinary nature.

Continuous computational problems arise in diverse areas of science and engineering. Examples include path and multivariate integration, approximation, optimization, as well as operator equations. Typically, only partial and/or noisy information is available, and the aim is to solve the problem within a given error tolerance using the minimal amount of computational resources. For example, in high-dimensional integration one wants to compute an ϵ-approximation to the integral with the minimal number of function evaluations. Here it is crucial to identify first the relevant variables of the function. Understanding the complexity of such problems and construction of efficient algorithms is both important and challenging. The current seminar attracted 35 participants from nine different countries all over the world. About 30 % of them were young researchers including PhD students. There were 25 presentations covering in particular the following topics:

- High-dimensional problems
- Tractability
- Computational stochastic processes
- Compressive sensing
- Random media
- Computational finance
- Noisy data
- Learning theory
- Biomedical learning problems
- Markov chains

There were three introductory talks to recent developments in PDE with random coefficients, learning theory and compressive sensing. A joint session with the Dagstuhl Seminar 15392 "Measuring the Complexity of Computational Content: Weihrauch Reducibility and Reverse Analysis" stimulated the transfer of ideas between the two different groups present in Dagstuhl.

The work of the attendants was supported by a variety of funding agencies. This includes the Deutsche Forschungsgemeinschaft, the Austrian Science Fund, the National Science Foundation (USA), and the Australian Research Council.

As always, the excellent working conditions and friendly atmosphere provided by the Dagstuhl team have led to a rich exchange of ideas as well as a number of new collaborations. Selected papers related to this seminar will be published in a special issue of the Journal of Complexity.

7.52 Measuring the Complexity of Computational Content: Weihrauch Reducibility and Reverse Analysis

Organizers: Vasco Brattka, Akitoshi Kawamura, Alberto Marcone, and Arno Pauly

Seminar No. 15392

Date: September 20–25, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.77

Participants: Vasco Brattka, Matthew de Brecht, Damir D. Dzhafarov, Fernando Ferreira, Willem L. Fouché, Cameron Freer, Guido Gherardi, Peter Hertling, Denis R. Hirschfeldt, Jeffry L. Hirst, Rupert Hölzl, Hajime Ishihara, Akitoshi Kawamura, Takayuki Kihara, Ulrich Kohlenbach, Alexander P. Kreuzer, Stéphane Le Roux, Alberto Marcone, Kenshi Miyabe, Antonio Montalbán, Carl Mummert, Eike Neumann, Paulo Oliva, Ludovic Patey, Arno Pauly, Matthias Schröder, Victor Selivanov, Paul Shafer, Dieter Spreen, Klaus Weihrauch, Keita Yokoyama, Kazuto Yoshimura, Martin Ziegler

Reducibilities such as many-one, Turing or polynomial-time reducibility have been an extraordinarily important tool in theoretical computer science from its very beginning. In recent years these reducibilities have been transferred to the continuous setting, where they allow to classify computational problems on real numbers and other (continuous) data types.

On the one hand, Klaus Weihrauch's school of computable analysis and several further researchers have studied a concept of reducibility that can be seen as an analogue of many-one reducibility for functions on such data. The resulting structure is a lattice that yields a refinement of the Borel hierarchy and embeds the Medvedev lattice. Theorems of for-all-exists form can be easily classified in this structure.

On the other hand, Stephen Cook and Akitoshi Kawamura have independently introduced a polynomial-time analogue of Weihrauch's reducibility, which has been used to classify the computational complexity of problems on real numbers and other objects. The resulting theory can be seen as a uniform version of the complexity theory on real numbers as developed by Ker-I Ko and Harvey Friedman.

The classification results obtained with Weihrauch reducibility are in striking correspondence to results in reverse mathematics. This field was initiated by Harvey Friedman and Stephen Simpson and its goal is to study which comprehension axioms are needed in order to prove certain theorems in second-order arithmetic. The results obtained so far indicate that Weihrauch reducibility leads to a finer uniform structure that is yet in basic agreement with the non-uniform results of reverse mathematics, despite some subtle differences.

Likewise one could expect relations between weak complexity theoretic versions of arithmetic as studied by Fernando Ferreira et al., on the one hand, and the polynomial-analogue of Weihrauch reducibility studied by Cook, Kawamura et al., on the other hand.

While the close relations between all these approaches are obvious, the exact situation has not yet been fully understood. One goal of our seminar was to bring researchers from the respective communities together in order to discuss the relations between these research topics and to create a common forum for future interactions.

We believe that this seminar has worked extraordinarily well. We had an inspiring meeting with many excellent presentations of hot new results and innovative work in progress, centred around the core topic of our seminar. In an Open Problem Session many challenging current research questions have been addressed and several of them have been solved either during the seminar or soon afterwards, which underlines the unusually productive atmosphere of this meeting.

A bibliography that we have compiled during the seminar witnesses the substantial amount of research that has already been completed on this hot new research topic up to today.

This report includes abstracts of many talks that were presented during the seminar, it includes a list of some of the open problems that were discussed, as well as the bibliography.

Altogether, this report reflects the extraordinary success of our seminar and we would like to use this opportunity to thank all participants for their valuable contributions and the Dagstuhl staff for their excellent support!

7.53 Circuits, Logic and Games

Organizers: Mikołaj Bojańczyk, Meena Mahajan, Thomas Schwentick, and Heribert Vollmer

Seminar No. 15401

Date: September 27 to October 2, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.105

Participants: Sreejith Ajithkumar, Christoph Berkholz, Olaf Beyersdorff, Mikołaj Bojańczyk, Michaël Cadilhac, Thomas Colcombet, Aiswarya Cyriac, Samir Datta, Anuj Dawar, Michael Elberfeld, Martin Grohe, Rohit Gurjar, Anselm Haak, Kristoffer Arnsfelt Hansen, Jan Johannsen, Juha Kontinen, Andreas Krebs, Klaus-Jörn Lange, Nutan Limaye, Kamal Lodaya, Meena Mahajan, Pierre McKenzie, Arne Meier, Stefan Mengel, David A. Mix Barrington, Anish Mukherjee, Anca Muscholl, Charles Paperman, Jean-Eric Pin, B. V. Raghavendra Rao, Thomas Schwentick, Luc Segoufin, Sebastian Siebertz, Karteek Sreenivasaiah, Howard Straubing, Dimitri Surinx, Thomas Thierauf, Jacobo Torán, Jan Van den Bussche, Jonni Virtema, Heribert Vollmer, Nils Vortmeier, Thomas Zeume

This seminar was the third in this series, the earlier two being Dagstuhl Seminars 06451 in November 2006 and 10061 in February 2010.

Goals of the Seminar

Over the years, there has been a lot of interplay between circuit complexity and logic. There are tight connections between small-depth circuit classes and fragments and extensions of first-order logic, and ideas from games and finite model theory have provided powerful lower bound techniques for circuits.

In recent years, there has been an impressive and sustained growth of interest and activity in the intersection of finite model theory and Boolean circuit complexity. The central aim of the seminar was to bring together researchers from these two areas to further strengthen the mutual fertilisation. Given the ubiquitousness of algebraic techniques in circuit complexity, the seminar also included arithmetic circuit complexity in its ambit.

The seminar focussed on the following specific topics:

- The algebraic approach to circuit complexity with its applications to finite model theory
- The logic-circuit connection, with a particular emphasis on circuit lower bounds that trigger results in finite model theory like separations between logics
- New connections between uniformity conditions on circuit families and logical predicates
- Structural complexity and circuit lower bounds inherently using methods from logic and algebra
- Proof systems with low circuit complexity
- Dynamic complexity: understanding the dynamic expressive power of small depth circuit classes

Organization of the Seminar and Activities

The seminar had the participation of 43 members from 11 countries.

The organisers attempted to create a schedule with a judicious mix of survey talks, focussed talks, and free time for unstructured discussions. Participants were invited to present their work and to communicate state-of-the-art advances. Since the participants came from diverse communities, the organisers invited some of them to give long survey-style talks in specific sub-areas. There were five such talks, listed below.

1. Olaf Beyersdorff. Lower bounds: from circuits to QBF proof systems.
   This talk surveyed the relatively new area of proof systems for establishing falseness of fully quantified Boolean formulas. It demonstrated techniques by which lower bounds in ciruit complexity can be tranferred to lower bounds on the sizes of such proofs.
2. Thomas Colcombet. Combinatorial Expressions and Lower Bounds.
   This talk described an elegant formalism, combinatorial expressions, that captures bounded depth circuits manipulating infinite data in specified restrictive ways, and showed how one may obtain indefinability results in this model.
3. Anuj Dawar. Lower Bounds for Symmetric Circuits.
   This talk described the recently formalised circuit model of symmetric circuits, its connections with logical definability, and a lower bound technique using games.
4. Martin Grohe. Color Refinement: A Simple Partitioning Algorithm with Applications From Graph Isomorphism Testing to Machine Learning.
   This talk described exciting connections between higher-dimensional generalisations of the extremely simple colour

refinement aglorithm and a linear programming approach to testing isomorphism.

5. Nutan Limaye. Arithmetic Circuit Lower Bounds.
   This talk surveyed the recent explosion of results concerning size lower bounds in restricted models of algebraic computation, using techniques which seem essentially combinatorial in nature.

In addition, 20 other participants gave short talks on some of their recent work relevant to the seminar theme. These talks covered results in two-variable first-order logic; dynamic complexity; graph colouring; database theory; circuit lower bounds; logics on words; and semigroup techniques. There was also a short session on Thursday devoted to discussing interesting open problems.

Concluding Remarks and Future Plans

The organizers regard the seminar as being quite successful. Most participants felt that they learnt new things from other areas, and were hopeful of using such ideas to make progress in their own research areas.

One aspect noted by the organizers was that a lot of the work discussed at the seminar used techniques from algebra. In fact, there was even a suggestion that if there is a future seminar in this series, it could be called "Circuits, Logic, and Algebra" instead of "Circuits, Logic, and Games".

The organizers are grateful to the Scientific Directorate of the Center for its support of this seminar.

7.54 Self-Assembly and Self-Organization in Computer Science and Biology

Organizers: Vincent Danos and Heinz Koeppl

Seminar No. 15402

Date: September 27 to October 2, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.9.125

Participants: Jacob Beal, Carsten Beta, Koan Briggs, José Antonio Carillo, Matthew Cook, Maria Rita D'Orsogna, Vincent Danos, Hanno Hildmann, Jonathan P. Hill, Heinz Koeppl, Nikita Kruk, Emmanuel Levy, Yuri Maistrenko, Vincent Noireaux, Matthew J. Patitz, John H. Reif, Matthias Sachs, Christian Scheideler, Rebecca Schulman, Friedrich Simmel, Adrian Sosic, Sandro Stucki, Petr Šulc, Serhiy Yanchuk

The Seminar brought together researchers from molecular biology, molecular modeling and theoretical computer scientists with interest in formal models of molecular computation and self-organization. Molecular biology provides a rich substrate to implement molecular computation and complex self-assembly algorithms. The Seminar featured several talks on DNA-assembly systems, that to-date represents the most advanced molecular substrate for self-assembly. The increase in the achieveable complexity of such molecular structures asks for a formal description and analysis of those systems using methods from theoretical computer science. The Seminar was successful in identifying common problem statements and in establishing a common scientific language. Apart from self-assembly, the broader term self-organization was mostly represented by research on swarming or self-propelled particle (SPP) models. The common feature of SPP systems and self-assembly is the emergence of global structures through local interaction rules (self-assembled structure vs swarms or flocks). One contribution also featured the combination of swarming and self-assembly system in terms of nucleation studies. Moreover, novel methodological overlap between simulation algorithms for molecular self-assembly and simulation algorithms for SPP systems were identified and elaborated during the workshop.

The seminar was structured as a regular workshop with morning and afternoon sessions but plenty of time was allocated for discussions after each talk. For the first such Dagstuhl Seminar no working groups were defined. For follow-up Seminars on the same topic we aim to additionally define working groups that may also deliver preliminary research results and initiation of new collaborations.

Although the workshop was very interdisciplinary we were able to arrange the presentations into sessions of a coherent theme. The feedback of participants was extremely positive, stating that they could really profit from the technical discussions that accompanied every presentation and that were performed in the free time. Correspondingly, several new collaborations across disciplines were initiated at the seminar.

7

7.55 Multimodal Manipulation Under Uncertainty

Organizers: Jan Peters, Justus Piater, Robert Platt, and Siddhartha Srinivasa
Seminar No. 15411

Date: October 4–9, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.10.1

Participants: Ron Alterovitz, Brenna D. Argall, Yasemin Bekiroglu, Kostas Bekris, Dmitry Berenson, Bastian Bischoff, Jeannette Bohg, Oliver Brock, Matei Ciocarlie, Fan Dai, Renaud Detry, Mehmet R. Dogar, Aaron M. Dollar, Roderic A. Grupen, Simon Hangl, David Hsu, Leslie Pack Kaelbling, Marek S. Kopicki, Dirk Kraft, Norbert Krüger, Ville Kyrki, Ales Leonardis, Shuai Li, Maxim Likhachev, Tomás Lozano-Pérez, Matthew T. Mason, Mark Moll, Duy Nguyen-Tuong, Erhan Öztop, Jan Peters, Justus Piater, Robert Platt, Maximo A. Roa, Veronica Santos, Siddhartha Srinivasa, Ales Ude, Francisco Valero Cuevas, Jeremy L. Wyatt, Michael Zillich

While robots have been used for decades to perform highly specialized tasks in engineered environments, robotic manipulation is still crude and clumsy in settings not specifically designed for robots. There is a huge gap between human and robot capabilities, including actuation, perception, and reasoning. However, recent developments such as low-cost manipulators and sensing technologies place the field in a good position to make progress on robot manipulation in unstructured environments. Various techniques are emerging for computing or inferring grasp configurations based on object identity, shape, or appearance, using simple grippers and robot hands.

Beyond grasping, a key ingredient of sophisticated manipulation is the management of *state information and its uncertainty*. One approach to handling uncertainty is to develop grasping and manipulation skills that are robust to environmental variation. Another approach is to develop methods of interacting with the environment in order to gain task-relevant information, for example, by touching, pushing, changing viewpoint, etc. Managing state information and uncertainty will require a tight combination of perception and planning. When the sensor evidence is unambiguous, the robot needs to be able to recognize that and perform the task accurately and efficiently. When greater uncertainty is present, the robot needs to adjust its actions so that they will succeed in the worst case or it needs to gain additional information in order to improve its situation. Different sensing modalities as well as world models can often be combined to good effect due to their complementary properties.

This seminar discussed research questions and agendas in order to accelerate progress towards robust manipulation under uncertainty, including topics such as the following:

- Is there a master algorithm or are there infinitely many algorithms that solve specialized problems? Can we decompose multimodal manipulation under uncertainty into I/O boxes? If so, what would these be?
- Do we prefer rare-feedback / strong-model or frequent-feedback / weak-model approaches? Is there a sweet spot in between? Is this the way to think about underactuated hands?
- What are useful perceptual representations for manipulation? What should be the relationship between perception and action? What kind of perception is required for reactive systems, planning systems, etc.?
- How do we do deformable-object manipulation? What planning methods, what types of models are appropriate?
- How should we be benchmarking manipulation? What kind of objects; what kind of tasks should be used?
- How should humans and robots collaborate on manipulation tasks? This question includes humans collaborating with autonomous robots as well as partially-autonomous robots acting under human command.

In the area of perception, we concluded that the design of representations remains a central issue. While it would be beneficial to develop representations that encompass multiple levels of abstraction in a coherent fashion, it is also clear that specific visual tasks suggest distinct visual representations.

How useful or limiting is the engineering approach of decomposing functionality into separate modules? Although this question was heavily debated, the majority view among seminar participants was that modules are useful to keep design complexity manageable for humans, and to keep the event horizon manageable for planning systems. It seems that to build more flexible and powerful systems, modules will need to be more strongly interconnected than they typically are these days. Fundamental challenges lie in the specification of each module and of their interconnections. There is a lot of room for creative innovation in this area.

Benchmarking questions were discussed chiefly in the context of the YCB Object Set. Specific benchmarks were suggested and discussed, covering perception and planning in the context of autonomous manipulation.

7.56 Dynamic Traffic Models in Transportation Science

Organizers: José R. Correa, Tobias Harks, Kai Nagel, Britta Peis, and Martin Skutella

Seminar No. 15412

Date: October 4–9, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.10.19

Participants: Umang Bhaskar, Roberto Cominetti, José R. Correa, Martin Gairing, Yannai A. Gonczarowski, Tobias Harks, Martin Hoefer, Max Klimm, Ekkehard Köhler, Felix König, Jannik Matuschke, Kai Nagel, Neil Olver, Britta Peis, Rahul Savani, Marco Scarsini, Miriam Schlöter, Daniel Schmand, Marc Schröder, Alexander Skopalik, Martin Skutella, Nicolás E. Stier-Moses, Martin Strehler, Chris M. J. Tampère, Veerle Timmermans, Laura Vargas-Koch, Bernhard von Stengel, Dave P. Watling

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

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

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

- Under which conditions do microscopic simulation models and dynamic traffic assignment models admit an equilibrium?
- Is an equilibrium efficiently (polynomial time) computable?
- Which models lead to multiple equilibria and how do the parameters of a learning process influence the resulting equilibrium outcome?
- What are the implications of possible intractability results (PPAD-hardness) on the plausibility of existing models?
- how do we compute optimal (or approximatively) network designs or traffic light controls subject to dynamic equilibrium constraints in polynomial time?

The seminar brought together leading researchers from three different communities – Simulations (SIM), Dynamic Traffic Assignment (DTA) and Algorithmic Game Theory (AGT) – and identified ways to narrow the existing gap between complex simulation based models and the existing theory. Among other points, the seminar initiated a systematic study of the complexity of equilibrium computations for DTA models – which is the core task when resolving dynamic traffic assignment problems. Equilibrium computation and its complexity status is a core topic in AGT. The seminar provided an excellent forum for a discourse of these questions between the DTA, SIM and AGT community which initiated several novel research questions and directions. The seminar also stimulated a conceptual discourse regarding the validity of DTA and microscopic simulation models in terms of their predictive power and use for optimization based approaches.

Overall, the seminar was a big success both in terms of stimulating new and very fruitful collaborations between so far separate communities and also with respect to novel insights and results on traffic equilibria and related concepts. We got enthusiastic feedback from many participants which is also reflected in the survey conducted by Dagstuhl.

7

7.57 Rack-Scale Computing

Organizers: Babak Falsafi, Tim Harris, Dushyanth Narayanan, and David A. Patterson

Seminar No. 15421

Date: October 11–16, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.10.35

Participants: Gustavo Alonso, Yungang Bao, Angelos Bilas, Peter Corbett, Paolo Costa, Christina Delimitrou, Felix Eberhardt, Lars Eggert, Babak Falsafi, Paolo Faraboschi, Christof Fetzer, Steve Furber, Jana Giceva, Matthew P. Grosvenor, Boris Grot, Hermann Härtig, Tim Harris, Maurice Herlihy, Matthias Hille, Torsten Hoefler, Konstantinos Katrinis, Kimberly Keeton, John Kim, Christoph M. Kirsch, Sergey Legtchenko, Martin Maas, Sue Moon, Andrew W. Moore, Dushyanth Narayanan, Jörg Nolte, Mark H. Oskin, Simon Peter, Andreas Polze, Danica Porobic, Zoran Radovic, Kaveh Razavi, Randolf Rotta, Ant Rowstron, Stefan Schmid, Bernhard Schräder, Malte Schwarzkopf, Liuba Shrira, Jens Teubner, Gael Thomas, Jana Traue, Leendert van Doorn, Haris Volos, Bernard Wong, Noa Zilberman, Ferad Zyulkyarov

Rack-scale computing is an emerging research area concerned with how we design and program the machines used in data centers. Typically, these data centers are built from racks of equipment, with each rack containing dozens of discrete machines connected by Ethernet or by InfiniBand. Over the last few years researchers have started to weaken the boundaries between these individual machines, leading to new "rack-scale" systems. These architectures are being driven by the need to increase density and connectivity between servers, while lowering cost and power consumption.

Initial commercial systems provide high-density processor nodes connected through an in-machine interconnect to storage devices or to external network interfaces (e.g., HPE Moonshot, or SeaMicro Fabric Compute). Many ideas are now being explored in research projects – e.g., the use of custom system-on-chip processors in place of commodity chips, the use of emerging non-volatile-memory technologies or stacked Flash in place of disks, and the use of silicon photonics and wireless links for communication within or between rack-scale systems. In addition, researchers are exploring how systems software, language runtime systems, and programming models can evolve for these new architectures.

This seminar sought to bring together researchers working on different parts of these problems. We structured the seminar around a small number of invited introductory talks accompanied by break-out sessions and a series of four poster sessions. The poster sessions permitted everyone to have an opportunity to present their own work (if they wished to), and enabled many parallel discussions to continue at the same time around different posters.

7.58 Genomic Privacy

Organizers: Jean Pierre Hubaux, Stefan Katzenbeisser, Bradley Malin, and Gene Tsudik

Seminar No. 15431

Date: October 18–23, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.10.5

Participants: Luk Arbuckle, Erman Ayday, Marina Blanton, Dan Bogdanov, Emiliano De Cristofaro, Zekeriya Erkin, Jacques Fellay, Kay Hamacher, Zhicong Huang, Jean Pierre Hubaux, Mathias Humbert, Aniket Kate, Stefan Katzenbeisser, Florian Kerschbaum, Oliver Kohlbacher, Florian Kohlmayer, Alexander Kaitai Liang, Huang Lin, Bradley Malin, Adam Molyneaux, Muhammad Naveed, Jun Pang, Fabian Prasser, Manuel Prinz, Jean-Louis Raisaro, Kurt Rohloff, Dominique Schröder, Vitaly Shmatikov, Sean Simmons, Adam Davison Smith, Thorsten Strufe, Qiang Tang, Carmela Troncoso, Juan Ramon Troncoso Pastoriza, Gene Tsudik, Paulo Jorge Veríssimo, Xiaofeng Wang

This report documents the program and the outcomes of Dagstuhl Seminar 15431 "Genomic Privacy". The current rise of personalized medicine is based on increasing affordability and availability of individual genome sequencing. Impressive recent advances in genome sequencing have ushered a variety of revolutionary applications in modern healthcare and epidemiology. In particular, better understanding of the human genome as well as its relationship to diseases and response to treatments promise improvements in preventive and personalized healthcare.

At the same time, human genetics has become a "big data" science. For roughly a decade, specific tests for Single Nucleotide Polymorphisms (SNPs), e.g., markers corresponding to specific diseases, have been well established. Furthermore, research in pharmaco-genomics, which currently relies on SNPs, has helped improve drug treatment for cancer and cardiac patients. The methodology of genotyping, which takes into account hundreds to thousands of variations in positions in the genome, has tremendously increased the amount of data acquired during diagnosis. Personalized genotyping has become commercially available from several sources (such as 23andMe). Full genome sequencing and genome-wide association studies are moving towards full deployment in clinical practice. In 2000, the cost of sequencing one human genome was US$2.5 billion. Today, the price of US$200 for genome sequencing is approaching reality. Considering the benefits for (public) health and potential cost savings, widespread acquisition, storage, and usage of personal genomes is guaranteed to happen soon.

However, because of the human genome's highly sensitive nature, this progress raises important privacy and ethical concerns, which simply cannot be ignored. A digitized genome represents one of the most sensitive types of human (personal) identification data. Even worse, a genome contains information about its owner's close relatives. Furthermore, correlations with individual data sets from so-called "omics-technologies" pose even bigger threats on privacy. Leakage of personal genomic information can lead a wide variety of attacks, many of which are not yet fully understood. Whether accidentally or intentionally revealed, a digitized genome cannot be revoked or modified. Consequently, secrecy of personal genomic data is of paramount importance. Furthermore, genomic data, unlike other types of highly sensitive information (even national secrets), does not lose its sensitivity over time. Even worse, the mechanisms available to interpret genomic data improve over time, which means that it is unclear at the moment how much sensitive information a genome encodes and which consequences a genomic data breach has. Furthermore, it is likely that genomic data will not only be used personally to support medical treatments; great promise lies in its use in large-scale genetic studies for personalized medicine as well as common ancestry and genetic compatibility tests. Therefore, simply encrypting genomic data at rest is not a viable option and new ways of protection need to be devised.

The second Dagstuhl Seminar on Genomic Pricacy concentrated on the following topics:

- Technical solutions for genomic privacy: the participants discussed technical solutions to enable genomic data privacy, even in the presence of untrusted computing environments, and investigated technical protection techniques that can be used for this purpose.
- Integration of genomic and physiological data: For medical purposes, genomic data often needs to be correlated with clinical and physiological data. For example, clinical studies may require finding correlations between physiological data reported during hospital stays and genomic information. So far, most technical solutions for the protection of genomic data focused on securely storing DNA data itself, but did not discuss the complex problem of combining it with physiological data.
- Protection of sensitive data within large-scale genome-wide

association studies: Although large-scale genomic studies offer many advantages for medical research, they pose many privacy problems. Most prior technical solutions focus on protection of a single human genome and do not scale multitudes of genomes. It remains a challenge to devise scalable techniques.

7.59 Duality in Computer Science

Organizers: Mai Gehrke, Achim Jung, Victor Selivanov, and Dieter Spreen

Seminar No. 15441

Date: October 25–30, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.10.66

Participants: Adriana Balan, Libor Barto, Ulrich Berger, Nick Bezhanishvili, Vasco Brattka, Matthew de Brecht, Max Dickmann, Abbas Edalat, Martin H. Escardo, Willem L. Fouché, Silvio Ghilardi, Jean Goubault-Larrecq, Georges Hansoul, Reinhold Heckmann, Achim Jung, Klaus Keimel, Dexter Kozen, Andreas Krebs, Clemens Kupke, Alexander Kurz, Jimmie D. Lawson, Matteo Mio, M. Andrew Moshier, Robert Myers, Dirk Pattinson, Daniela Petrisan, Michael Pinsker, Luca Reggio, Giuseppe Rosolini, Matthias Schröder, Peter M. Schuster, Niels Schwartz, Victor Selivanov, Vladimir Shavrukov, Alex Simpson, Dieter Spreen, Paul Taylor, Marcus Tressl, Samuel J. van Gool, Yde Venema, Steven J. Vickers, Klaus Weihrauch

Aims of the seminar

Duality allows one to move between an algebraic world of properties and a spacial world of individuals and their dynamics, thereby leading to a change of perspective that may, and often does, lead to new insights. Because computer science is fundamentally concerned both with specification of programs and the dynamics of their executions, dualities have given rise to active research in a number of areas of theoretical computer science. In this seminar we particularly wanted to concentrate on applications of duality in semantics for continuous data with special focus on probability in computation, algebra and coalgebra, and applications in complexity theory.

The seminar

Our call for participation was exceptionally successful and right up to the actual start of the meeting we were in danger of exceeding the number of places allocated. We see this as a vindication of our aim of bringing these researchers together for exchanging ideas centred around the common topic of duality. The talks offered fell quite naturally into groupings which allowed us to adopt a fairly thematic programme structure:

Day 1, morning session: Duality and classical algebra. Talks by Libor Barto, Michael Pinsker, Max Dickmann, and Marcus Tressl.

Day 1, afternoon session: Duality and categories. Talks by Paul Taylor, Steve Vickers, and Pino Rosolini.

Day 2, morning session: Duality and topology. Talks by Matthew de Brecht, Mathias Schröder, Reinhold Heckmann, and Jean Goubault-Larrecq.

Day 2, afternoon session: Alternative views on duality. Talks by Niels Schwartz, George Hansoul, Rob Myers, and Alexander Kurz.

Day 3, morning session: Duality and coalgebra. Talks by Adriana Balan, Dirk Pattinson, Ulrich Berger, and Samuel J. van Gool.

Day 4, morning session: Duality and domain theory. Talks by Jimmie Lawson, Abbas Edalat, Achim Jung, and Klaus Keimel.

Day 4, afternoon session: Duality and logic. Talks by Peter Schuster, Martín Escardó, Vladimir Shavrukov, and Vasco Brattka.

Day 5, morning session: Duality and probability. Talks by Willem Fouché, Dexter Kozen, Daniela Petrişan, and Drew Moshier.

Final thoughts

As always, Dagstuhl staff were incredibly efficient and helpful which allowed all of us, including the organisers, to focus on the exchange of ideas and plans for joint work. We are sincerely grateful to them for their hospitality and professionalism.

Fig. 7.18
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of (above) TreeState Productions GmbH, www.treestate.de and (below) Dr. Christian Lindig.

7.60 Approaches and Applications of Inductive Programming

Organizers: José Hernández-Orallo, Stephen H. Muggleton, Ute Schmid, and Benjamin Zorn

Seminar No. 15442

Date: October 25–30, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.10.89

Participants: Umair Zafrulla Ahmed, Tarek R. Besold, Harold Boley, Andrew Cropper, Luc De Raedt, Cesar Ferri Ramirez, Sumit Gulwani, José Hernández-Orallo, Petra Hofstedt, Frank Jäkel, Susumu Katayama, William B. Langdon, Fernando Martinez-Plumed, Martin Möhrmann, Stephen H. Muggleton, Hila Peleg, Ruzica Piskac, Maria José Ramirez Quintana, Ute Schmid, Michael Siebers, Rishabh Singh, Gustavo Soares, Armando Solar-Lezama, Claes Strannegård, Lorijn van Rooijen, Janis Voigtländer, Christina Zeller, Benjamin Zorn

Inductive programming research addresses the problem of learning (mostly declarative) programs from incomplete specifications, such as input/output examples, observed traces, or constraints. Beginning in the 1960s, this area of research was initiated in artificial intelligence (AI) exploring the complex intellectual cognitive processes involved in producing program code which satisfies some specification. Furthermore, applications of AI for software engineering are investigated resulting in methodologies and techniques for automating parts of the program development process. Inductive programming can be seen as a very special subdomain of machine learning where the hypothesis space consists of classes of computer programs.

Nowadays, researchers working on inductive programming are distributed over different communities, especially inductive logic programming, evolutionary programming, grammar inference, functional programming, and programming languages and verification. Furthermore, similar approaches are of interest in programming by demonstration applications for end-user programming as well as in cognitive models of inductive learning.

The recent release of FlashFill as a plug-in inductive programming tool for Microsoft Excel is an impressive demonstration that inductive programming research has matured in such a way that commercial applications become feasible. Similarly, the field has attracted widespread interest in computer science as a whole, as illustrated by the recent review article published by the Communications of the ACM [1].

In the seminar, we brought together researchers from different areas of computer science – especially from machine learning, AI, declarative programming, and software engineering – and researchers from cognitive psychology interested in inductive learning as well as in teaching and learning computer programming. Furthermore, participants from industry presented current as well as visionary applications for inductive programming.

We addressed many aspects which partially were identified as relevant topics during the previous Dagstuhl Deminar 13502 (http://www.dagstuhl.de/13502). In particular, we had the following sessions for presentations:

- Session: General techniques, languages and systems for inductive logic and inductive functional programming.
- Session: End-user programming, programming by example and applications
- Session: Program synthesis and transformation
- Cognitive Aspects of Induction

In addition, we had several systems demos and tutorials (some of them 'hands-on'):

- System demo on Metagol.
- Hands-on-Tutorial: The MagicHaskeller Library and Server System
- Tutorial: FlashMeta SDK for creating programming-by-example tools
- Tutorial: Sketch synthesis infrastructure

The seminar also included a DemoFest, where several systems were demonstrated in small groups in a relaxed atmosphere.

The first and second days the following topics were identified and further discussed in working groups during the rest of the seminar:

- Benchmarks, Evaluation, and Applications
- General-Purpose IP Infrastructures and Applicability and Evaluation Criteria for IP Approaches in the Context of AI
- Probabilities in IP

Concluding remarks and future plans

In the final panel discussion the results of the seminar were summarised as well as future plans.

Regarding the seminar, there were several suggestions that

topics should be more mixed, instead of grouping them too much into "silo" sessions. About the format of the seminar, there was a general agreement that the change from half a week to one week had been beneficial, and that the DemoFest had been a real success. Indeed, the possibility of having several independent demos earlier in the week and more demo sessions (or DemoFests) was a possible suggestion for subsequent meetings.

The following topics were elaborated about future actions:

- Make the community and the area more visible through tutorials and workshops at major conferences, or summer schools.
- Integrate tools, demos, videos, tutorials and other kinds of material at www.inductive-programming.org/resources.html
- Focus on benchmarks and a common representation language for problems, and use the inductive-programming website to publish the benchmarks. Organise competitions (or hackathons building apps that use the underlying IP engines).
- Revitalise the mailing list (at the moment of writing this report the list is fully operative again, see http://www.inductive-programming.org/)
- Attract people from other areas (e.g., cognitive robotics).
- Change the frequency of the meetings to a 1-year cadence, with perhaps Dagstuhl every other year and a competition or summer school in between.

Overall, the main conclusion can be summarised as the realisation of a very significant progress in techniques and its exploitation in new applications, so it is now time to strengthen the visibility of the IP research and its community, for which this Dagstuhl seminar has served as a lever.

References

1 S. Gulwani, J. Hernández-Orallo, E. Kitzelmann, S. H. Muggleton, U. Schmid, and B. Zorn. *Inductive programming meets the real world*. Communications of the ACM, 58, 11, 2015.

7.61 Verification of Evolving Graph Structures

Organizers: Parosh Aziz Abdulla, Fabio Gadducci, Barbara König, and Viktor Vafeiadis

Seminar No. 15451

Date: November 1–6, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.1

Participants: Mohamed-Faouzi Atig, Parosh Aziz Abdulla, Peter Backes, Paolo Baldan, Ahmed Bouajjani, Andrea Corradini, Aiswarya Cyriac, Giorgio Delzanno, Cezara Dragoi, Constantin Enea, Javier Esparza, Fabio Gadducci, Silvio Ghilardi, Holger Giese, Christoph Haase, Annegret Habel, Reiko Heckel, Alexander Heußner, Lukas Holik, David Janin, Christina Jansen, Bengt Jonsson, Joost-Pieter Katoen, Barbara König, Tomer Kotek, Narayan Kumar Krishnan, Leen Lambers, Michele Loreti, Roland Meyer, Thomas Noll, Fernando Orejas, Eugenio Orlandelli, Oded Padon, Detlef Plump, Chris Poskitt, Arend Rensink, Ahmed Rezine, Leila Ribeiro, Xavier Rival, Arnaud Sangnier, Richard Trefler, Viktor Vafeiadis, Tomas Vojnar, Thomas Wies, Florian Zuleger

Despite significant progress in recent years, verification still remains a challenging task for hardware and software systems. A particularly complex verification problem is the analysis of graph-like structures that may modify their topology during runtime. The main reason for the difficulty is that some features give rise to infinite state spaces. Examples include variables ranging over unbounded domains, timing constraints, dynamic process creation, heap manipulation, multi-threading, and dynamically allocated data structures. An additional source of complication is that the underlying graphs may be continuously evolving. There is no a priori bound on the size of the graphs that may arise when modelling the run of a program, and the graph shapes may change during a given execution.

This challenge has prompted several successful lines of research, developing novel techniques such as shape analysis, separation logic, forest automata, and several graph transformation-based approaches. Although specialized tools have been developed in each application area, a considerable amount of effort is needed to develop uniform frameworks that yield efficient yet general solutions.

This seminar brought together researchers interested in developing precise and scalable techniques for the analysis of graph manipulations, i.e., techniques that are able to handle the challenges that arise in current verification problems. These challenges require novel developments and the combination of techniques from a wide range of different areas including model checking and dynamic and static program analysis. By creating collaboration opportunities we hope to substantially increase the size of the systems that can be tackled and the precision of analysis that can be achieved.

Hence the main goal of this seminar was to enhance common understanding and cross-fertilization, highlighting connections among the approaches via tutorials and working groups, with the explicit purpose to enhance interaction. Discussion topics included:

- the definition of uniform frameworks in which to integrate methods for graph analysis that have been proposed by the different research communities;
- the development of new abstraction techniques for pushing the state-of-the-art of graph algorithms in program verification and model checking applications; and
- the identification of research areas in which the analysis of graph manipulation may play an important role, such as the analysis of security protocols, social networks, adaptive networks, and biological systems.

We invited four representatives of the different communities to give tutorial talks in order to introduce fundamental concepts and techniques. Specifically, the following four tutorial talks took place on the first day of the seminar:

- Tomas Vojnar: Shape Analysis via Symbolic Memory Graphs and Its Application for Conversion of Pointer Programs to Container Programs
- Giorgio Delzanno: Graphs in Infinite-State Model-Checking
- Arend Rensink: Verification Techniques for Graph Rewriting
- Viktor Vafeiadis: Separation Logic

On Tuesday and Thursday we organized the following working groups in order to discuss more specific topics which were of interest to a substantial part of the participants:

- Benchmarks and Application Domains
- Specification Languages for Graphs
- Ownership
- Graph Rewriting for Verification

The organizers would like to thank all the participants and speakers for their inspiring talks and many interesting discussions. Furthermore we would like to acknowledge Christina Jansen and Eugenio Orlandelli who helped to write and prepare the full report. A special thanks goes to the Dagstuhl staff who were a great help in organizing this seminar.

7

7.62 Artifact Evaluation for Publications

Organizers: Bruce R. Childers, Grigori Fursin, Shriram Krishnamurthi, and Andreas Zeller

Seminar No. 15452

Date: November 1–4, 2015 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.5.11.29

Participants: Bruce R. Childers, Neil Chue Hong, Tom Crick, Jack W. Davidson, Camil Demetrescu, Roberto Di Cosmo, Jens Dittrich, Dror Feitelson, Sebastian Fischmeister, Grigori Fursin, Ashish Gehani, Matthias Hauswirth, Marc Herbstritt, David R. Kaeli, Shriram Krishnamurthi, Anton Lokhmotov, Martin Potthast, Lutz Prechelt, Petr Tuma, Michael Wagner, Andreas Zeller

Computer systems researchers have developed numerous artifacts that encompass a broad collection of software tools, benchmarks, and data sets. These artifacts are used to prototype innovations, evaluate trade-offs and analyze implications. Unfortunately, methods used in the evaluation of computing system innovation are often at odds with sound science and engineering practice. The ever-increasing pressure to publish more and more results poses an impediment to *accountability*, which is a key component of the scientific and engineering process. Experimental results are not usually disseminated with sufficient metadata (i.e., software extensions, data sets, benchmarks, test cases, scripts, parameters, etc.) to achieve repeatability and/or reproducibility. Without this information, issues surrounding trust, fairness and building on and comparing with previous ideas becomes problematic. Efforts in various computer systems research sub-communities, including programming languages/compilers, computer architecture, and high-performance computing, are underway to address the challenge.

This Dagstuhl Perspectives Workshop brought together stakeholders of associated CSR sub-communities to determine synergies and to identify the most promising directions and mechanisms to push the broader community toward accountability. The Perspectives Workshop assessed current efforts, shared what does and doesn't work, identified additional processes, and determined possible incentives and mechanisms. The outcomes from the workshop, including recommendations to catalyze the community, are separately documented in an associated Dagstuhl Manifesto.

7.63 Vision for Autonomous Vehicles and Probes

Organizers: André Bruhn, Atsushi Imiya, Aleš Leonardis and Tomas Pajdla

Seminar No. 15461

Date: November 8–13, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.36

Participants: José M. Alvarez, Juan Andrade-Cetto, Steven S. Beauchemin, Florian Becker, Sven Behnke, Johannes Berger, Andrés Bruhn, Darius Burschka, Daniel Cremers, Krzysztof Czarnecki, Cédric Demonceaux, Michael Felsberg, Friedrich Fraundorfer, Yasutaka Furukawa, Rafael Garcia, Antonios Gasteratos, Andreas Geiger, Michal Havlena, Heiko Hirschmuller, Ben Huber, Atsushi Imiya, Reinhard Koch, Takashi Kubota, Lazaros Nalpantidis, Mikael Persson, Thomas Pock, Danil V. Prokhorov, Sebastian Ramos, Hayko Riemenschneider, Torsten Sattler, Davide Scaramuzza, Bernt Schiele, Jürgen Sturm, Niko Sünderhauf, Akihiko Torii, Raquel Urtasun, Vladyslav Usenko, David Vázquez Bermudez, Andreas Wendel, Christian Winkens

Computer vision plays a key role in advanced driver assistance systems (ADAS) as well as in exploratory and service robotics. Visual odometry, trajectory planning for Mars exploratory rovers and the recognition of scientific targets in images are examples of successful applications. In addition, new computer vision theory focuses on supporting autonomous driving and navigation as applications to unmanned aerial vehicles (UAVs) and underwater robots. From the viewpoint of geometrical methods for autonomous driving, navigation and exploration, the on-board calibration of multiple cameras, simultaneous localisation and mapping (SLAM) in non-human-made environments and the processing of non-classical features are some of current problems. Furthermore, the adaptation of algorithms to long image sequences, image pairs with large displacements and image sequences with changing illumination is desired for robust navigation and exploration. Moreover, the extraction of non-verbal and graphical information from environments to remote driver assistance is required.

Based on these wide range of theoretical interests from computer vision for new possibility of practical applications of computer vision and robotics, 38 participants (excluding organisers) attended from variety of countries: 4 from Australia, 3 from Austria, 3 from Canada, 1 from Denmark, 11 from Germany, 1 from Greece, 1 from France, 3 from Japan, 4 from Spain, 2 from Sweden, 4 from Switzerland and 3 from the US.

The seminar was workshop style. The talks are 40 mins and 30 mins for young researchers and for presenters in special sessions. The talks have been separated into sessions on aerial vehicle vision, under water and space vision, map building, three-dimensional scene and motion understanding as well as a dedicated session on robotics. In these tasks, various types of autonomous systems such as autonomous aerial vehicles, under water robots, field and space probes for remote exploration and autonomous driving cars were presented. Moreover, applications of state-of-the-art computer vision techniques such as global optimization methods, deep learning approaches as well as geometrical methods for scene reconstruction and understanding were discussed. Finally, with Seminar 15462 a joint session on autonomous driving was organised.

The working groups are focused on "Sensing," "Interpretation and Map building" and "Deep leaning." Sensing requires fundamental methodologies in computer vision. Low-level sensing is a traditional problem in computer vision. For applications of computer-vision algorithms to autonomous vehicles and probes, reformulation of problems for various conditions are required. Map building is a growing area including applications to autonomous robotics and urban computer vision. Today, application to autonomous map generation involves classical SLAM and large-scale reconstruction from indoor to urban sizes. Furthermore, for SLAM on-board and on–line computation is required. Deep learning, which goes back its origin to '70s, is a fundamental tool for image pattern recognition and classification. Although the method showed significant progress in image pattern recognition and discrimination, for applications to spatial recognition and three-dimensional scene understanding, we need detailed discussion and developments.

Through talks-and-discussion and working-group discussion, the seminar clarified that for designing of platforms for visual interpretation and understanding of three-dimensional world around the system, machine vision provides fundamental and essential methodologies. There is the other methodology which uses computer vision as a sensing system for the acquisition of geometrical data and analysis of motion around cars. For these visual servo systems, computer vision is a part of the platform for intelligent visual servo system. The former methodology is a promising one to provide a fundamental platform which is common to both autonomous vehicles, which are desired for consumer intelligence, and probes, which are used for remote exploration.

7.64 The Mobile Revolution – Machine Intelligence for Autonomous Vehicles

Organizers: Wolfram Burgard, Uwe Franke, Markus Enzweiler, and Mohan Trivedi

Seminar No. 15462

Date: November 10–13, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.62

Participants: Michael Aeberhard, Jan Becker, Klaus Bengler, Claus Brenner, Wolfram Burgard, Erik Coelingh, Michael Darms, Markus Enzweiler, Ryan Eustice, Uwe Franke, Dariu M. Gavrila, Alex Goldberg, Ralf G. Herrtwich, Ulrich Hofmann, Michael James, Serge Lambermont, Antonio M. Lápez Pena, Chris Mansley, Markus Maurer, Karsten Mühlmann, Urs Muller, Michel Parent, Mikael Persson, Raul Rojas, Torsten Sattler, Steven E. Shladover, Christoph Stiller, Matthias Strauß, Mohan Trivedi, Sadayuki Tsugawa, D. Scott Williamson, Hermann Winner, Hans-Joachim Wünsche

Motivation and Perspective

Machine intelligence, robotics and computer vision, formerly rather peripheral disciplines of computer science, are in fact already with us today and have a familiar embodiment – the modern vehicle. Systems that are currently available strongly couple interdisciplinary fundamental research with complex practical realizations. The vision of autonomous vehicles in particular has a surprisingly long history with first prototypical implementations going back to the early 1980s. What started then as a dream of pioneers such as Ernst Dickmanns is actually happening right now – we are on the verge of a mobile revolution with self-driving vehicles as its central foundation. The tremendous progress made in the last years has been sparked by the increased methodical and technically availability of better sensors, sophisticated algorithms, faster computers and more data.

But, we are not quite there yet. Autonomous systems make extreme demands on system performance, quality, availability, reliability and verification that significantly increase with the rising degree of automation. Such diverse requirements give rise to numerous problems and open questions that are currently addressed in substantial academic and industrial research activities in many fields of computer science and engineering. Extraordinarily positive innovation effects result from the knowledge transfer between industry and academia, as successfully demonstrated by initiatives such as Uni-DAS or DRIVE-U. The increasing relevance and interest in the computer science community, particularly in the fields of robotics, computer vision and machine learning is evident through an abundance of papers and workshops at major computer science conferences.

This seminar has brought together the leading experts from both academia and industry to discuss the state-of-the-art, identify further research directions and refine the overall vision of intelligent autonomous vehicles into a consistent and practicable picture.

Seminar Topics and Structure

The ultimate design goal for autonomous systems is to mimic human behavior in terms of understanding and effortlessly acting within a dynamic human-inhabited environment. Although artificial sensors emulating the human sensory systems are nowadays widely available, current autonomous systems are still far behind humans in terms of understanding and acting in real-world environments. The chief reason is the (theoretical and practical) unavailability of methods to reliably perform perception, recognition, understanding and action on a broad scale, i.e. not limited to isolated problems.

Following the classical perception-action cycle, the central topics of the seminar have evolved around four key questions posed from the perspective of an autonomous vehicle:

- What do I perceive and how can I interpret this?
- Where am I and what do I do next?
- How can I build up experience and learn?
- Am I capable of this task?

More specifically, the seminar has stimulated research and discussions through several talks on the following topics:

- Intelligent Robotics
- Digital Maps
- Human-centered Intelligent Vehicles
- Verification and Validation
- Limitations and Perspectives

The seminar was held in a very interactive workshop style allowing for ample time for thorough discussions. There were four main sessions with talks and discussions, c.f. the seminar schedule in Section 4 of the full report, focusing on autonomous driving projects, mapping and localization, sensing, as well as evaluation and approval. The first session on state-of-the-art autonomous driving projects has been co-organized with Seminar 15461 as a joint session.

7.65 Symbolic Computation and Satisfiability Checking

Organizers: Erika Ábrahám, Pascal Fontaine, Thomas Sturm, and Dongming Wang

Seminar No. 15471

Date: November 15–20, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.71

Participants: John Abbott, Erika Ábrahám, Bernd Becker, Martin Bromberger, Christopher W. Brown, Shaowei Cai, Florian Corzilius, James H. Davenport, Pascal Fontaine, Stephen Forrest, Jürgen Gerhard, Maximilian Jaroschek, Dejan Jovanovic, Tim A. King, Konstantin Korovin, Marek Kosta, Laura Kovács, Gereon Kremer, Wolfgang Küchlin, Viktor Levandovskyy, Klaus Meer, David Monniaux, Chenqi Mou, Mizuhito Ogawa, Andrew Joseph Reynolds, Yosuke Sato, Karsten Scheibler, Tobias Schubert, Viorica Sofronie-Stokkermans, Thomas Sturm, Laurent Voisin, Christoph M. Wintersteiger, Patrick Wischnewski, Kazuhiro Yokoyama

The seminar focused on satisfiability checking for combinations of first-order logic and subclasses thereof with arithmetic theories in a very liberal sense, also covering quantifiers and parameters.

The development of decision procedures for corresponding theories started in the early 20th century in the area of mathematical logic. In the second half of the 20th century it played a prominent role within the development of algebraic model theory. Finally, around 1970, one important research line, viz. algebraic decision methods for real arithmetics, shifted its focus from theoretical results towards practically feasible procedures. That research line was one of the origins of an area known today as *symbolic computation* or *computer algebra*.

More recently, the *satisfiability checking* community, which originated from propositional SAT solving and which is surprisingly disconnected from symbolic computation, began to develop highly interesting results with a particular focus on existential decision problems, following the track of SAT solving towards industrial applications. Powerful *satisfiability modulo theories (SMT)* solvers were developed, which enrich propositional SAT solving with components for different theories. We understand satisfiability checking in a broad sense, covering besides SMT solving also *theorem proving* with arithmetic.

The two communities of *symbolic computation* and *satisfiability checking* have been quite disjoint, despite strong reasons for them to discuss together. The communities share interests, e.g., examining arithmetic expressions, that are central to both. As a matter of fact, the symbolic computation community has been mostly unaware of basic insights in the satisfiability checking community, such as the efficiency of conflict-driven search with learning, as well as of their fundamental requirements, e.g., incrementality or explanations in the unsatisfiable case. Vice versa, researchers in satisfiability checking have adopted decision procedures from symbolic computation, such as CAD for real closed field, only quite naively, so that they do not really benefit from the considerable experience gained by the original community during 45 years. It is our hope that our seminar contribute to bringing the two communities together, and that they will be much stronger at tackling problems that currently defeat them both, separately.

The seminar offered its participants an opportunity to exchange knowledge about existing methods and applications, to push forward the communication of needs and interests, and to draw attention to challenging open research questions. The participants included researchers from all relevant research areas and with affiliations in academia and as well as in industry. The program was a balanced combination of presentations and tutorials, but also offering time for small group discussions and exchange of ideas.

To the best of our knowledge, the seminar was the first global meeting of the two communities of symbolic computation and satisfiability checking. We are confident that it will initiate cross-fertilization of both fields and bring improvements for both satisfiability checking and symbolic computation, and for their applications.

7

7.66 Programming with "Big Code"

Organizers: William W. Cohen, Charles Sutton, and Martin T. Vechev

Seminar No. 15472

Date: November 15–18, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.90

Participants: Miltos Allamanis, Earl Barr, Jason Breck, Swarat Chaudhuri, William W. Cohen, Premkumar T. Devanbu, Shi Han, Kenneth Heafield, Abram Hindle, Suresh Jagannathan, Christopher M. Jermaine, Dongsun Kim, Dana Movshovitz-Attias, Tien N. Nguyen, Sebastian Proksch, Christopher Quirk, Veselin Raychev, Armando Solar-Lezama, Charles Sutton, Daniel Tarlow, Martin T. Vechev, Nicolas Voirol, Eran Yahav, Andreas Zeller, Xin Zhang

The main objective of the seminar was to bring together several research communities which have so far been working separately on the emerging topic of "Big Code" and to foster a new community around the topic. Over the last 4–5 years there have been several developments and interesting results involving "Big Code" all spanning a wide range of fields and conferences: the seminar brought these communities together and enabled them to interact for the first time.

The program was structured as a series of talks interspersed with discussion. Almost all of seminar participants gave a talk on their latest research. Even though the initial plan was to include special discussion sessions, each talk triggered so much discussion, both during the talk itself, and also after, that there was no need for specific discussion slots. We believe the seminar was successful in setting the right atmosphere for open ended discussion and obtained the desired affect of triggering much organic interaction.

Only the last day (morning) included a short wrap-up discussion session focusing on the future of the area, defining common data sets and future challenges the community can address. That discussion is summarized in the working group report (as part of the full report).

The seminar was highly inter-disciplinary involving experts from programming languages, software engineering, machine learning and natural language processing. Further, it brought together research groups from Europe, Asia and U.S., all working on the topic of "Big Code", and raised awareness and familiarity with what different research groups are working on.

The talks and discussions spanned several topics including: the kinds of statistical methods used (e.g., n-gram models, recurrent neural networks, graphical models, probabilistic grammars, etc), new programming applications that can benefit from these models (e.g., code completion, code search, code similarity, translating natural language to code, etc), and the interaction between these. Some of the presentations were more of an introductory/overview nature while others focused on the more technical aspects of particular programming tools and machine learning models.

After two days of presentations and discussions, we used the last day of the seminar (before lunch) to summarize the discussions and to outline a future research direction. A suggestion enthusiastically embraced by everyone was to create a web site which lists the current data sets, challenges, tools and research groups working on the topic. The view was that this will not only enable existing groups to compare their tools on common problems and data sets but will also make it much easier for other research groups and graduate students to get into the area and to start contributing. It also serves as a useful instrument for raising awareness about the topic:

We have now created this web site and have made it available here: http://learnbigcode.github.io/.

In a short time, several groups have started contributing by uploading links to tools, data sets and challenges.

Overall, the seminar was successful both in terms of stimulating new and fruitful interaction between research communities that were working in the area but were separated so far, but also in setting a common agenda moving forward. Due to the high interest and feedback from this seminar, we anticipate that in a year or two from now, we will be ready to propose a larger seminar on the topic.

7.67 Evaluation in the Crowd: Crowdsourcing and Human-Centred Experiments

Organizers: Daniel Archambault, Tobias Hoßfeld, and Helen C. Purchase

Seminar No. 15481

Date: November 22–27, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.103

Participants: Daniel Archambault, Benjamin Bach, Kathrin Ballweg, Rita Borgo, Alessandro Bozzon, Sheelagh Carpendale, Remco Chang, Min Chen, Stephan Diehl, Darren J. Edwards, Sebastian Egger, Sara Fabrikant, Brian D. Fisher, Ujwal Gadiraju, Neha Gupta, Matthias Hirth, Tobias Hoßfeld, Jason Jacques, Radu Jianu, Christian Keimel, Andreas Kerren, Stephen G. Kobourov, Bongshin Lee, David Martin, Andrea Mauri, Fintan McGee, Luana Micallef, Sebastian Möller, Babak Naderi, Martin Nöllenburg, Helen C. Purchase, Judith Redi, Peter Rodgers, Dietmar Saupe, Ognjen Scekic, Paolo Simonetto, Tatiana von Landesberger, Ina Wechsung, Michael Wybrow, Michelle X. Zhou

In various areas of computer science like visualization, graphics, or multimedia, it is often required to involve the users, e.g. to measure the performance of the system with respect to users, e.g. to measure the user perceived quality or usability of a system. A popular and scientifically rigorous method for assessing this performance or subjective quality is through formal experimentation, where participants are asked to perform tasks on visual representations and their performance is measured quantitatively (often through response time and errors). For the evaluation of the user perceived quality, users are conducting some experiments with the system under investigation or are completing user surveys. Also in other scientific areas like psychology, such subjective tests and user surveys are required. One approach is to conduct such empirical evaluations in the laboratory, often with the experimenter present, allowing for the controlled collection of quantitative and qualitative data. Crowdsourcing platforms can address these limitations by providing an infrastructure for the deployment of experiments and the collection of data over diverse user populations and often allows for hundreds, sometimes even thousands, of participants to be run in parallel over one or two weeks. However, when running experiments on this platform, it is hard to ensure that participants are actively engaging with the experiment and experimental controls are difficult to implement. Often, qualitative data is difficult, if not impossible, to collect as the experimenter is not present in the room to conduct an exit survey. Finally, and most importantly, the ethics behind running such experiments require further consideration. When we post a job on a crowdsourcing platform, it is often easy to forget that people are completing the job for us on the other side of the machine.

The focus of this Dagstuhl seminar was to discuss experiences and methodological considerations when using crowdsourcing platforms to run human-centred experiments to test the effectiveness of visual representations in these fields. We primarily target members of the human-computer interaction, visualization, and applied perception research as these communities often engage in human-centred experimental methodologies to evaluate their developed technologies and have deployed such technologies on crowdsourcing platforms in the past. Also, we engaged researchers that study the technology that makes crowdsourcing possible. Finally, researchers from psychology, social science and computer science that study the crowdsourcing community participated and brought another perspective on this topic. In total, 40 researchers from 13 different countries participated in the seminar. The seminar was held over one week, and included topic talks, stimulus talks and flash ('late breaking') talks. In a 'madness' session, all participants introduced themselves in a fast-paced session within 1 minutes. The participants stated their areas of interest, their expectations from the seminar, and their view on crowdsourcing science. The major interests of the participants were focused in different working groups:

- Technology to support Crowdsourcing
- Crowdworkers and the Crowdsourcing Community
- Crowdsourcing experiments vs laboratory experiments
- The use of Crowdsourcing in Psychology research
- The use of Crowdsourcing in Visualisation research
- Using Crowdsoursing to assess Quality of Experience

The abstracts from the different talks, as well as the summary of the working groups can be found in the full report. Apart from the report, we will produce an edited volume of articles that will become a primer text on (1) the crowdsourcing technology and methodology, (2) a comparison between crowdsourcing and lab experiments, (3) the use of crowdsourcing for visualization, psychology, and applied perception empirical studies, and (4) the nature of crowdworkers and their work, their motivation and demographic background, as well as the relationships among people forming the crowdsourcing community.

Fig. 7.19
Impressions from the 25th anniversary celebratory colloquium of Schloss Dagstuhl, July 3rd, 2015. Photos courtesy of (above) TreeState Productions GmbH, www.treestate.de and (below) Dr. Christian Lindig.

7.68 Social Concepts in Self-organising Systems

Organizers: Ada Diaconescu, Stephen Marsh, Jeremy Pitt, Wolfgang Reif, and Jan-Philipp Steghöfer

Seminar No. 15482

Date: November 22–27, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.127

Participants: Gerrit Anders, Kirstie Bellman, Nelly Bencomo, Olivier Boissier, Jean Botev, Pompeu Casanovas, Ada Diaconescu, Babak Esfandiari, Sebastian Götz, Hanno Hildmann, Jan Kantert, Peter R. Lewis, Stephen Marsh, Christian Müller-Schloer, Gauthier Picard, Jeremy Pitt, Wolfgang Reif, Ingo Scholtes, Jan-Philipp Steghöfer, Sven Tomforde, Leon van der Torre, Laurent Vercouter

There are two exciting trends in computing that motivated this seminar. On the one hand, large-scale self-organising systems gain traction in real-world settings, e.g., in the autonomous control of the power grid or in personal transportation scenarios. On the other hand, our lives are more and more pervaded by socio-technical systems that rely on the interaction of existing, complex social systems and technical systems that in many ways mirror and form the social relationships of their users. The seminar brought together researchers from a variety of domains to discuss the technical, legal, and social issues these trends incur. One focus was how social concepts can be formalised and implemented to make technical self-organising systems more robust and efficient. The other focus was how technology shapes the social system and vice versa.

Use of Social Concepts in Self-Organising and Socio-Technical Systems

The seminar's first focus is motivated by the requirements of large-scale self-organising systems. The more such systems have to take their environment into account, the more open, and the more heterogeneous they are, the more important social concepts become [1]. If a population of agents is no longer developed, deployed, maintained, and controlled by a single company or institution, the goals of the agents no longer concur – especially, the individual sub-goals no longer necessarily imply the overall system goal. In such cases, social constructs can help encourage cooperation between the agents. The presence of norms as an explicit expression of acceptable behaviour [3], of a trust management system to encourage reciprocity [4], or of a form of computational justice to settle disputes within the system [5] are measures that have been discussed in the scientific community for these ends. In this way, the systems form a legal reality that establishes certain rules and regulations within the system. This legal reality must be in accordance with the legal system under whose jurisdiction these systems work.

The second focus is how technical systems interact with and influence existing social systems. With the increasing dependence of society on computation and on complex artificial systems, their influence on human-computer interaction, and on inter-human interaction becomes a topic of concern (see, e.g., [2]). One aspect of this is the novel challenge of managing an online identity, made necessary by the representations of human users in technical systems that are, necessarily, an abstraction of the real user. Another aspect is the increasing reliance of human users on these technical aids and the potential of negative effects on the users accompanied with this. Such effects can range from infringement of privacy, to withholding of relevant information, and even to targeted manipulation.

Results of the Seminar

The seminar was highly interactive, with a lot of time dedicated to plenum discussions. Talks were used as impulses to stimulate discussion and working groups focused on particular aspects that the participants deemed particularly important.

A number of talks addressed the implications of using social concepts in technical systems from different angles and featured insights into existing technical solutions, e.g., for computational justice, trust, and ethical behaviour, as well as the observable effects of these solutions. Likewise, incentives and how social constructs influence them was a recurring theme. The discussions following these talks addressed important issues such as the relationship between system and user values, goals and the rules designed to achieve these goals, possible attacks on socio-technical systems, quantifiable incentives, and self-determination in technical systems.

A different set of talks was aimed at understanding the way social systems and technology interact, e.g., how the social organisation of the human users is represented in the technical system and becomes evident in the interactions in that system. An overview of the interplay of legal and technical systems was also provided, with important insights into the connection between technical feasibility and legal admissibility. This set of talks encouraged discussion geared towards governance, power, and the representation of values in technical systems, as well as on how to represent existing social systems in technical systems and how both the social system and its representation evolve over time.

Based on the discussions and the input from the talks, three working groups were formed that focused on discussing different aspects of the use of social concepts in self-organising systems. Their main aims and contributions are as follows:

Understanding: A first step was to consider what the notion of "social" means in the context of the technical systems regarded in the seminar. Based on a brief literature survey, the working group determined that social means that an organisation exists, that the welfare of the individuals and the organisation is regarded, and that the relations between individuals and between individuals and the organisation are a concern. A second step was to stipulate that formalising social values leads to the individuals behaving in a way that recognises their social obligations and responsibilities. Finally, the notion of "socially-sensitive design" was introduced to denote that both the design process and the system itself must be socially sensitive.

Engineering: The main concerns were how to make different social concepts usable in technical systems, how to select the fitting social construct for a specific problem, how to measure its effectiveness, and how to combine several social concepts. The working group suggested a pattern language to express selection criteria, implementation approaches, and consequences, as well as a set of metrics that make it possible to evaluate the impact of the social concept in the technical system.

Dynamics: The discussion developed towards how the social and technical components of socio-technical systems interact with each other and how the resulting dynamic aspects influence these systems. A total of six challenges were identified, the most important of which pertains to how the interaction between different social concepts that provides "checks and balances" in social systems can be transferred to technical systems. Further problems that were discussed are conflict resolution and power distribution, as well as the influence of technical systems on society and where the responsibility for this influence lies.

Future Work

The seminar participants agreed that the topic is timely and relevant and that there are a number of open issues that need to be addressed in the future. Possible venues for future elaboration of these issues are the SASOˆST workshops[47], held annually at the IEEE Conference on Self-Adaptive and Self-Organising Systems, as well as a number of other projects currently in discussion. In particular, the organisers are discussing ways to provide an overview of the state of the art of the field as well as a research roadmap and opportunities to specifically discuss the impact self-organising and socio-technical systems will have on society.

References

1 L. Vercouter and G. Muller. L.i.a.r.: Achieving social control in open and decentralized multiagent systems. *Applied Artificial Intelligence*, 24(8):723–768, 2010.

2 J. Botev. Anonymity, immediacy and electoral delegation in socio-technical networked computer systems. In M. Horbach, editor, *Informatik 2013, 43. Jahrestagung der Gesellschaft für Informatik e.V. (GI), Informatik angepasst an Mensch, Organisation und Umwelt, 16.-20. September 2013, Koblenz*, volume 220 of *LNI*, page 1164. GI, 2013.

3 T. Balke, C. da Costa Pereira, F. Dignum, E. Lorini, A. Rotolo, W. Vasconcelos, and S. Villata. Norms in MAS: Definitions and related concepts. In G. Andrighetto, G. Governatori, P. Noriega, and L. van der Torre, editors, *Normative Multi-Agent Systems*, volume 4 of *Dagstuhl Follow-Ups*, pages 1–31. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013.

4 L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model of trust and reputation. In *Proc. of the 35th Hawaii International Conference on System Sciences (HICSS'02)*, pages 188–196. IEEE Computer Society Press, 2002.

5 J. Pitt, D. Busquets, and R. Riveret. Formal models of social processes: The pursuit of computational justice in self-organising multi-agent systems. In *7th IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2013, Philadelphia, PA, USA, September 9-13, 2013*, pages 269–270. IEEE, 2013.

[47] http://sasost.isse.de/

7.69 Approximate and Probabilistic Computing: Design, Coding, Verification

Organizers: Antonio Filieri, Marta Kwiatkowska, Sasa Misailovic, and Todd Mytkowicz

Seminar No. 15491

Date: November 29 to December 4, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.151

Participants: Sara Achour, David Bindel, Mateus Araújo Borges, James Bornholt, Luca Bortolussi, Tomas Brázdil, Andreas Peter Burg, Luis Ceze, Eva Darulova, Alessandra Di Pierro, Luis María Ferrer Fioriti, Antonio Filieri, Jaco Geldenhuys, Andreas Gerstlauer, Lars Grunske, Jane Hillston, Ulya R. Karpuzcu, Joost-Pieter Katoen, Marta Kwiatkowska, Rupak Majumdar, Dimitrios Milios, Sasa Misailovic, Subhasish Mitra, Todd Mytkowicz, Ravi Nair, Karthik Pattabiraman, Adrian Sampson, Karin Strauss, Willem Visser, Herbert Wiklicky

Uncertainty and approximation are becoming first class concepts in software design and development. Many application domains, including biology, multimedia processing, finance, engineering, and social sciences, need software to formalize and study intrinsically uncertain phenomena. Furthermore, the ubiquity of software, especially driven by the Internet and mobility – such as driving applications that estimate routes, speech processing applications that estimate most likely sentences, or fitness applications that estimate heart-rate – require software engineers to design their applications taking into account unpredictable and volatile operational conditions, and noisy data, despite the limited support provided by current unintuitive design and quality assurance methodologies. Finally, the hardware community is designing devices that trade result accuracy for computational efficiency and energy saving, providing only probabilistic guarantees on the correctness of the computed results.

Several research communities are independently investigating methodologies and techniques to model, analyze, and manage uncertainty in and through software systems. These areas include (1) probabilistic model checking, (2) quantitative software analysis, (3) probabilistic programming, and (4) approximate computing. However, despite the substantial overlap of interests, researchers from different communities rarely have the opportunity to meet at conferences typically tailored to single specific areas. Therefore, we organized this seminar as a forum for industrial and academic researchers from these areas to share their recent ideas, identify the main research challenges and future directions, and explore collaborative research opportunities on problems that span across the boundaries of the individual areas.

This report presents a review of each of the main areas covered by the seminar and summarizes the discussions and conclusions of the participants.

Acknowledgements. The organizers would like to express their gratitude to the participants and the Schloss Dagstuhl team for a productive and exciting seminar. We thank Prof. Martin Rinard for his support and contribution to the organization of the seminar. We thank Sara Achour for her help with preparing the full report.

7.70 Computational Metabolomics

Organizers: Sebastian Böcker, Juho Rousu, and Emma Schymanski

Seminar No. 15492

Date: November 29 to December 4, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.11.180

Participants: Felicity Allen, Nuno Bandeira, Sebastian Böcker, Corey Broeckling, Céline Brouard, Jacques Corbeil, Pieter Dorrestein, Kai Dührkop, P. Lee Ferguson, Franziska Hufsky, Gabi Kastenmüller, Tobias Kind, Oliver Kohlbacher, Daniel Krug, Kris Morreel, Steffen Neumann, Tomas Pluskal, Lars Ridder, Simon Rogers, Juho Rousu, Emma Schymanski, Huibin Shen, Christoph Steinbeck, Michael Stravs, Ales Svatos, Tom Wenseleers, Rohan Williams, David Wishart, Michael Anton Witting, Gert Wohlgemuth, Nicola Zamboni

Metabolomics has been referred to as the apogee of the omics-sciences, as it is closest to the biological phenotype. Mass spectrometry is the predominant analytical technique for detecting and identifying metabolites and other small molecules in high-throughput experiments. Huge technological advances in mass spectrometers and experimental workflows during the last decades enable novel investigations of biological systems on the metabolite level. But these advances also resulted in a tremendous increase of both amount and complexity of the experimental data, such that the data processing and identification of the detected metabolites form the largest bottlenecks in high throughput analysis. Unlike proteomics, where close co-operations between experimental and computational scientists have been established over the last decade, such cooperation is still in its infancy for metabolomics.

The Dagstuhl Seminar on Computational Metabolomics brought together leading experimental and computational side experts in a dynamically-organized seminar designed to foster the exchange of expertise. Overview talks were followed by breakout sessions on topics covering the whole experimental-computational continuum in mass spectrometry.

7.71 The Graph Isomorphism Problem

Organizers: László Babai, Anuj Dawar, Pascal Schweitzer, and Jacobo Torán

Seminar No. 15511

Date: December 13–18, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.12.1

Participants: Albert Atserias, László Babai, Christoph Berkholz, Jannis Bulian, Sourav Chakraborty, Maurice Chandoo, Bireswar Das, Anuj Dawar, Michael Elberfeld, Martin Fürer, Martin Grohe, Rohit Gurjar, Neil Immerman, Tommi Junttila, Sandra Kiefer, Johannes Köbler, Sebastian Kuhnert, Piyush P. Kurur, Daniel Lokshtanov, José Luis Lopez Presa, Eugene M. Luks, Brendan McKay, Daniel Neuen, Prajakta Nimbhorkar, Luis F. Núñez Chiroque, Yota Otachi, Michał Pilipczuk, Adolfo Piperno, Ilia Ponomarenko, Gaurav Rattan, David J. Rosenbaum, Alexander Russell, Patrick Scharpfenecker, Uwe Schöning, Pascal Schweitzer, Xiaorui Sun, Thomas Thierauf, Jacobo Torán, Yadu Vasudev, Oleg Verbitsky, Pengming Wang, John Wilmes, James B. Wilson, Angela Wu

The Graph Isomorphism Problem remains one of the two unresolved computational problems from Garey and Johnson's list dating back to 1979 of problems with unknown complexity status. In very rough terms the problems asks to decide whether two given graphs are structurally different or one is just a perturbed variant of the other. The problem naturally arises when one is faced with the task of classifying relational structures (e.g., chemical molecules, websites and links, road networks).

While the Graph Isomorphism Problem was intensively studied from the point of view of computational complexity in the 1980s and early 1990s, in later years progress became slow and interest in the problem stalled. However, recent years have seen the emergence of a variety of results related to graph isomorphism in a number of research areas including algorithmic group theory, finite model theory, combinatorial optimization and parameterized algorithms, not to mention graph theory itself. Indeed, having been open and quite prominent for such a long time, the Graph Isomorphism Problem is repeatedly attacked with the abundance of algorithmic techniques that have been developed over the decades. While this has not led to resolution of the problem, it has led to applications of methods originally developed for the Graph Isomorphism Problem in other areas (such as machine learning and constraint satisfaction problem solving). It has also sparked fascinating concepts in complexity theory, led to a thriving compilation of techniques in algorithmic group theory, the development of software packages (such as canonical labeling tools) and perpetuating effects in algorithmic graph theory in general.

While a lot of other computational problems have a specific community associated with them, resulting in dedicated conferences, the situation for the isomorphism problem is different. This is due to the fact that the background of people working on the isomorphism problem is quite diverse which leads to infrequent encounters at regular conferences or other events. Moreover, there is a big gap between theory and practice, a phenomenon verbalized by Brendan McKay as two distinct galaxies with very few stars in between them. Indeed, the algorithms that are asymptotically fastest in theory are very different to the ones that prove to be the fastest in practical implementations. The original motivation of the seminar was to bring together researchers working on the many topics closely related to the Isomorphism Problem to foster their collaboration.

However, the face of the seminar was to change, as one of the organizers (László Babai) published a proof on the arXiv (http://arxiv.org/abs/1512.03547) on the night before the seminar that shows that graph isomorphism can be solved in quasi-polynomial time (see the abstract to the talk below). This is the first improvement over the moderately exponential algorithm for general graphs by Luks from 1983. Babai gave three intense blackboard presentations each with a duration of two hours on the new quasi-polynomial time algorithm. Apart from the presentations, there were a number of excellent talks including expository surveys on recent advances in a variety of aspects of the Graph Isomorphism Problem as detailed below.

Overall a memorable event, we hope that the seminar has encouraged future collaboration across the different areas which eventually brings us closer to the theoretical and practical resolution of the problem.

7

7.72 Debating Technologies

Organizers: Iryna Gurevych, Eduard H. Hovy, Noam Slonim, and Benno Stein

Seminar No. 15512

Date: December 13–18, 2015 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.5.12.18

Participants: Khalid Al-Khatib, Jens Allwood, Carlos Alzate, Wolf-Tilo Balke, Yonatan Bilu, Elena Cabrio, Claire Cardie, Walter Daelemans, Ido Dagan, Anette Frank, Norbert Fuhr, Iryna Gurevych, Ivan Habernal, Graeme Hirst, Yufang Hou, Eduard H. Hovy, Christoph Lofi, Marie-Francine Moens, Brian Plüss, Vinodkumar Prabhakaran, Chris Reed, Nils Reiter, Ruty Rinott, Hinrich Schütze, Noam Slonim, Christian Stab, Manfred Stede, Benno Stein, Simone Teufel, Anita De Waard, Henning Wachsmuth

Why do people in all societies argue, discuss, and debate? Apparently, we do so not only to convince others of our own opinions, but because we want to explore the differences between our own understanding and the conceptualizations of others, and learn from them. Being one of the primary intellectual activities of the human mind, debating naturally involves a wide range of conceptual capabilities and activities, ones that have only in part been studied from a computational perspective in fields like computational linguistics and natural language processing. As a result, computational technologies supporting human debating are scarce, and typically still in their infancy. Recent decades, however, have seen the emergence and flourishing of many related and requisite computational tasks, including sentiment analysis, opinion and argumentation mining, natural language generation, text summarization, dialogue systems, recommendation systems, question answering, emotion recognition/generation, automated reasoning, and expressive text to speech.

This Dagstuhl seminar was the first of its kind. It laid the groundwork for a new interdisciplinary research community centered around debating technologies – computational technologies developed directly to enhance, support, and engage with human debating. The seminar brought together leading researchers from relevant communities to discuss the future of debating technologies in a holistic manner.

The seminar was held between 13 and 18 December 2015, with 31 participants from 22 different institutions. The event's sixteen sessions included 34 talks, thirteen themed discussions, three system demonstrations, and a hands-on "unshared" task. Besides the plenary presentations and discussions, the program included several break-out sessions and mock debates with smaller working groups. The presentations addressed a variety of topics, from high-level overviews of rhetoric, argument structure, and argument mining to low-level treatments of specific issues in textual entailment, argumentation analysis, and debating-oriented information retrieval. Collective discussions were arranged for most of these topics, as well as on more forward-thinking themes, such as the potential and limitations of debating technologies, identification of further relevant research communities, and plans for a future interdisciplinary research agenda.

A significant result of the seminar was the decision to use the term computational argumentation to put the community's various perspectives (argument mining, argument generation, debating technologies, etc.) under the same umbrella. By analogy with "computational linguistics", "computational argumentation" denotes the application of computational methods for analyzing and synthesizing argumentation and human debate. We identified a number of key research questions in computational argumentation, namely:

- How important are semantics and reasoning for real-world argumentation?
- To what extent should computational argumentation concern itself with the three classical rhetorical appeals of ethos (appeal to authority), pathos (appeal to emotion), and logos (appeal to reason)? Is it sufficient to deal with logos, or is there some benefit in studying or modelling ethos and pathos as well?
- What are the best ways of dealing with implicit knowledge?

A number of discussion questions at the seminar followed from these points, particularly in relation to the data and knowledge sources required for implementing and evaluating computational argumentation systems. For example, are currently available datasets sufficient for large-scale processing or for cross-language and cross-domain adaptation? Can we reliably annotate logos, ethos, and pathos? In any case, what sort of data would be considered "good" for a shared task in computational argumentation? Is it possible for computational argumentation to repeat the recent successes of "deep" natural language processing

by employing shallow methods on large masses of data? How does cultural background impact human argumentation, and is this something that computational models need to account for? Finding the answers to these and other questions is now on the agenda for our burgeoning research community.

8 Öffentlichkeitsarbeit
Public Relations and Outreach

8.1 Pressemitteilungen und Medienarbeit

8.1 *Press Releases and Media Work*

Die regelmäßige Erstellung und Herausgabe von Pressemitteilungen dient der verständlichen Verbreitung von aktuellen Informatikthemen. Die Darstellung des Konzepts von Schloss Dagstuhl kann dabei ebenfalls berücksichtigt werden. Pressemitteilungen und Berichterstattungen in diversen Medien – soweit bekannt – sind über das Internetportal von Schloss Dagstuhl[48] abrufbar.

Durch Unterstützung des Saarländischen Rundfunks steht Schloss Dagstuhl ein professionelles Reporterset zur Verfügung, welches Rundfunkjournalisten erlaubt, mit Seminarteilnehmern Interviews in digitaler verlustfreier Audioqualität zu führen.

Schloss Dagstuhl hat sich im Allgemeinen zur Anlaufstelle für Journalisten etabliert, die über bestimmte Informatikthemen, aber auch über Schloss Dagstuhl berichten möchten.

Schloss Dagstuhl verbreitet Neuigkeiten rund um sein Programm über soziale Netzwerkdienste wie Twitter und LinkedIn. Über Twitter-Nutzer @dagstuhl werden primär Programmankündigungen an aktuell ca. 940 Abonnenten verbreitet. Zunehmend nutzen aber auch Seminarteilnehmer den Dienst, um ihre Eindrücke vom Seminar mitzuteilen. Darüber hinaus werden über den Twitter-Nutzer @dblp_org Informationen über die Bibliographiedatenbank dblp an über 260 Abonnenten verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ unterhalten (derzeit etwa 640 Mitgliedern), 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 informatics 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].

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.

Schloss Dagstuhl has become a port of call for journalists seeking to report on specific informatics topics and/or on Schloss Dagstuhl itself.

News on the program of Schloss Dagstuhl are disseminated via social networks such as Twitter and LinkedIn. The Twitter handle @dagstuhl is used primarily to disseminate program announcements to about 940 followers, but is increasingly used also by Dagstuhl Seminar participants to share their impressions. Additionally, information about the dblp computer science bibliography is sent using the Twitter account @dblp_org, having more than 260 followers. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with about 640 members), which supports the networking of participants in Dagstuhl Seminars. Interesting news items pertaining to Schloss Dagstuhl are also disseminated. Additionally, interesting news about Schloss Dagstuhl are announced there.

8.2 Fortbildung

8.2 *Educational Training*

Schloss Dagstuhl engagiert sich im schulischen Bereich durch die Organisation einer jährlichen Lehrerfortbildung, die sich an Informatik- und Mathematiklehrer der gymnasialen Oberstufe im Saarland und in Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut Rheinland-Pfalz (PL) organisiert. Diese beiden Institute unterstützen die Fortbildung auch finanziell, indem sie die Kosten der Referenten tragen.

Jede Lehrerfortbildung dauert drei Tage; an jedem Tag werden in jeweils 3-stündigen Vorträgen zwei Informatikthemen vorgestellt. Die intensive Fortbildung richtet sich zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch können seit 2011 bis zu fünf Lehrer aus anderen Bundesländern teilnehmen.

Das Interesse an dieser Fortbildung stieg seit dem Beginn in 1991 stetig an und die 25. Lehrerfortbildung in

Schloss Dagstuhl holds an annual teacher training workshop specifically designed for teachers of upper secondary students working in the Saarland or the Rhineland Palatinate. The workshop is organized together with the Landesinstitut Pädagogik und Medien (LPM), Saarland, and the Pädagogisches Landesinstitut Rheinland-Pfalz (PL). These two institutes support the event also financially by assuming the costs of speakers.

Each workshops lasts three days; each day two computer science topics are presented in a three hour presentation each. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, since 2011 up to five teachers of other federal states can participate.

Interest in the workshop has risen steadily since the program began in 1991 and the 25th annual Dagstuhl Teacher Training Workshop, held at Schloss Dagstuhl on December 9–11, 2015, attracted more participants than

[48] http://www.dagstuhl.de/de/ueber-dagstuhl/presse/

Informatik, die vom 9. bis 11. Dezember 2015 statt fand, führte mehr Teilnehmer zusammen als jemals zuvor. Grund hierfür war nicht nur das hervorragende Programm, sondern auch das Kolloquium am Mittwoch, den 9. Dezember 2015, das Schloss Dagstuhl anlässlich der 25. Fortbildung ausgerichtet hat:

Einer kurzen Begrüßung durch den wissenschaftlichen Direktor folgten Grußworte von der saarländischen Ministerpräsidenten, Frau Annegret Kramp-Karrenbauer, und der rheinland-pfälzischen Ministerin für Bildung, Wissenschaft, Weiterbildung und Kultur, Frau Vera Reiß. Danach gaben Reinhard Wilhelm und Raimund Seidel einen Überblick über die vergangenen und einen Ausblick über zukünftige Lehrerfortbildungen in Dagstuhl. Das Kolloquium endete mit einem kleinen Sektempfang in dem Speisesaal von Schloss Dagstuhl. Insgesamt war die Feier mit ca. 80 Teilnehmern ein großer Erfolg.

Mehr Informationen zur Veranstaltung 2015 gibt es auf der Webseite der Veranstaltung[49].

25 Jahre Schloss Dagstuhl

8.3 *25th Anniversary of Schloss Dagstuhl*

Am 3. Juli 2015 feierte Schloss Dagstuhl sein 25. Bestehen mit einem Kolloquium für Gremienmitglieder und Dagstuhl-Freunde. Etwa 70 Personen haben an der Feier bei schönstem Wetter teilgenommen. Das Kolloquium startete mit einem wissenschaftlichen Programm mit Shriram Krishnamurthi, Thomas Lengauer und Wolfgang Thomas. Danach gab es Festreden von Matthias Kleiner, Präsident der Leibniz-Gemeinschaft, und Politikvertretern. Zu den Politikvertretern gehörten Jürgen Lennartz, Staatssekretär und Chef der Staatskanzlei des Saarlandes, und Thomas Deufel (Staatssekretär im Ministerium für Bildung, Wissenschaft, Weiterbildung und Kultur in Rheinland-Pfalz). Raimund Seidel und Reinhard Wilhelm rundeten das Programm mit einem kurzen Überblick über die Vergangenheit und Zukunft von Schloss Dagstuhl ab. Nach den Vorträgen und Reden wurde die Feier im Garten von Schloss Dagstuhl fortgeführt, die bis in die Nacht dauerte. Insgesamt war die Feier ein großer Erfolg und bei allen Gästen geschätzt.

ever before. This was not only the excellent program, but also the colloquium on Wednesday, December 9, 2015, Schloss Dagstuhl Castle has aligned on the occasion of the 25th event:

It started with greetings of Saarland's prime minister, Annegret Kramp-Karrenbauer, and Vera Reiß, minister of the ministry of education, science, and culture of Rhineland-Palatinate. Afterwards, Reinhard Wilhelm and Raimund Seidel gave an overview on the past and an outlook of future teacher education at Dagstuhl. This program was followed by a short reception. Altogether, the celebration with approximately 80 participants was a great success.

Details on the training in 2015 are available at the event webpage[49].

Schloss Dagstuhl celebrated its 25th anniversary on Friday, July 3, 2015, with a colloquium for board members and Dagstuhl friends. Approximately 70 participants attended the celebration in finest weather. The colloquium started in the afternoon with a scientific program by Shriram Krishnamurthi, Thomas Lengauer, and Wolfgang Thomas. Afterwards, the greetings of the president of the Leibniz Association, Matthias Kleiner, and representatives of the politics, Jürgen Lennartz (state secretary and head of the chancellery office of the Saarland), and Thomas Deufel (state secretary at the ministry of education, science, and culture of Rhineland-Palatinate) took place. Raimund Seidel and Reinhard Wilhelm completed the program by giving a brief overview on the past and the future of Dagstuhl. This program was followed by a garden party, which lasted well into the night. Altogether, the celebration was a great success and appreciated by all guests.

[49] http://www.dagstuhl.de/15503

9 Einrichtung und Service
Facilities and Services

Tagungsräume 9.1 *Conference Facilities*

Schloss Dagstuhl bietet drei Hörsäle für 25 bis 60 Personen. Alle Hörsäle sind mit einem Beamer, einem MS-Windows-Laptop und einer Audioanlage einschließlich Mikrophonen ausgestattet. Durch diese Technik werden Vorträge, Präsentationen und Live-Vorführungen auch verteilter Systeme optimal unterstützt. Mittels eines Presenter können Vortragende ihre vorbereiteten Materialien präsentieren, ohne zum Laptop oder Arbeitsplatz zurückkehren zu müssen.

Neben den Hörsälen gibt es im Zentrum sechs Seminarräume. Davon sind zwei mit modernen Beamern ausgestattet, während in einem ein großes Plasmadisplay montiert ist. Fünf Beamer auf Rollwagen stehen zusätzlich zur flexiblen Benutzung in allen Räumen zur Verfügung.

Die beiden größten Hörsäle sind jeweils mit mehreren Tafeln ausgestattet, während in den anderen Tagungsräumen jeweils große Whiteboards an den Wänden montiert sind. In einem Seminarraum kann durch eine spezielle Wandfarbe sogar eine ganze Wand als Whiteboard (über $12\mathrm{m}^2$) benutzt werden.

Daneben gibt es über das ganze Zentrum verteilt weitere Räume, in denen Gäste sich in entspannter Atmosphäre treffen und diskutieren können. Insbesondere am Abend zieht es viele Gäste in den Weinkeller und die Cafeteria, zwei der gemütlichsten Räume im Haus und hervorragend geeignet für die Fortsetzung einer produktiven Diskussion in angenehmer Atmosphäre.

Schloss Dagstuhl has three lecture halls with a seating capacity of 25 to 60 each. All lecture halls are equipped with a projector, an MS-Windows notebook, and an audio system including a microphone. These facilities not only enable talks and papers to be presented in an optimal manner but also permit online demonstrations of active and distributed systems to be given to large audiences. A presenter for use of those who wish to go through their presentations without physical access to a computer is also available.

In addition to the lecture halls, the center has six meeting rooms. Two are equipped with up-to-date projectors and one has a large plasma display at the wall. Five mobile projectors are available for use in all of the rooms.

Whereas the two main lecture halls are equipped with several blackboards, whiteboards are provided in the other rooms. One of the conference rooms features a complete "whiteboard wall" painted with a special paint which allows to use this whole wall (over $12\mathrm{m}^2$) as one large whiteboard.

The center also offers a spectrum of other spaces where guests can sit and work together in a relaxed atmosphere. In particular in the evening, guests gravitate towards the wine cellar and upstairs cafe, two of the coziest places in the house and great places for continuing with a productive discussion in a comfortable atmosphere.

Computer und Vernetzung 9.2 *Computers and Networks*

Schloss Dagstuhl ist über das DFN (Deutsches Forschungsnetz) mit zwei 100 Mbit/s Leitungen mit dem Internet verbunden. Während im Regelbetrieb eine Leitung für die Gäste reserviert ist, wird die andere für die Services von Schloss Dagstuhl sowie für die Mitarbeiter der verschiedenen Standorte benutzt. Jederzeit, z. B. im Fall eines Ausfalls, kann aber der gesamte Netzwerkverkehr über eine Leitung abgewickelt werden. Die Ausfallsicherung wird noch dadurch erhöht, dass die Endpunkte der Glasfaserleitungen nach Kaiserslautern und Saarbrücken in zwei verschiedenen Gebäuden liegen. Nahezu im ganzen Zentrum können sich Gäste über WLAN (IEEE 802.11 b, g, n) mit dem Internet verbinden. Der Zugriff erfolgt entweder über eduroam oder über eine Dagstuhl-eigene Kennung.

Neben dem Zugang über mitgebrachte Laptops, Tablet Computer oder Smartphones steht den Gästen ein Rechnerraum mit sechs Arbeitsplätzen zur Verfügung. Davon sind zwei Arbeitsplätze mit Apple Macs ausgerüstet, drei sind dedizierte MS-Windows-Arbeitsplätze und ein weiterer Arbeitsplatz mit Ubuntu Linux. Zusätzlich stehen den Gästen für die komfortable Arbeit an ihrem eigenen Rechner drei Arbeitspläze bestehend aus großem Monitor, externer Tastatur, Maus und Ethernetkabel zur Verfügung. Weitere Ethernet-Anschlüsse stehen ebenso zur Verfügung, um das WLAN wegen Bandbreite oder aus Kapazitätsgründen zu

Schloss Dagstuhl is connected to the Internet by the DFN (German Research Network) using two redundant 100 Mbit/s lines. Normally, one of the lines is used exclusively for our guests while the other one is reserved for Dagstuhl's Internet services and for the employees at the different locations. At any time, e.g. in case of failure, the whole traffic can be routed over one of the lines. The fail safe is increased by the fact that the endpoints of the two lines to Kaiserslautern and Saarbrücken are distributed over two separated buildings. Throughout the grounds guests have Internet access by Wi-Fi (IEEE 802.11 b, g, n). Access is either granted via eduroam or via a Dagstuhl-hosted private account.

Most of our guests prefer to use their laptops, tablet computers and smartphones, but Schloss Dagstuhl offers also one computer room with six workstations. Among them there are two Apple Macs, three dedicated MS-Windows workstations and one workstation providing Ubuntu Linux. For comfortably using their notebooks, Dagstuhl provides three work places consisting of a large display, an external keyboard, a mouse, and an Ethernet connection. More Ethernet cables are also provided to bypass the rate-restricted Wi-Fi connection. Two iPads, and upon request a MacBook Pro and a laptop with MS-Windows are also available for use throughout the grounds.

umgehen. Weiterhin bietet Schloss Dagstuhl seinen Gästen zwei iPads sowie auf Nachfrage ein MacBook Pro und einen Laptop mit MS-Windows.

Im Zentrum steht den Gästen ein Multifunktions-Farbdrucker mit Scanner und Kopierer zur Verfügung. Der Zugriff erfolgt über eine Weboberfläche, die das direkte Drucken zahlreicher Dokumentenformate erlaubt. Zu der IT-Ausstattung gehören weiterhin fünf Recherche-Arbeitsplätze inklusive Drucker in der Bibliothek sowie drei Laptops in den Seminarräumen.

Die tägliche Datensicherungen von Schloss Dagstuhl wird auf jeweils zwei unterschiedlichen Tape-Libraries gespeichert. Diese befinden sich in unterschiedlichen Gebäuden, so dass im Katastrophenfall (Brand, Überschwemmung, etc.) eine komplette Datensicherung zur Verfügung steht. Eine Library steht im Serverraum im sogenannten Neubau, die andere befindet sich im Technikraum des neu erbauten Gästehauses. Beide Gebäude sind über eine FiberChannel-Verbindung verbunden, die bei einer Entfernung von ca. 200 m eine Geschwindigkeit von 4 Gbit/s erreicht.

Schloss Dagstuhl provides a multifunction color printer with scanner and copier. Using a web front guests can upload and print the most used document formats without converting them. The center's IT equipment also includes five workstations including printers in the library for literature research, as well as three notebooks in the lecture halls.

The daily data backups of Schloss Dagstuhl are stored on two different tape libraries. These are located in different buildings to ensure that a complete backup is available in the event of a disaster (fire, flood, etc.). One of the libraries is located in the server room of the so-called New Building, while the other one is housed in the technical utilities room of the new Guest House. Both buildings are connected via a fiber-optic which reaches a speed of 4 Gbit/s for the 200 m distance.

Internet Dienste 9.3 *Internet Services*

Schloss Dagstuhl bietet allen Organisatoren und Gästen eine wachsende Anzahl webbasierter Dienste. Während der Vorbereitungsphase können alle Organisatoren tagesaktuell überprüfen, welche eingeladen Gäste bereits zu- oder abgesagt haben. Sie können ebenfalls einen (vorläufigen) Zeitplan auf der seminarspezifischen Webseite hochladen. Alle Teilnehmer können Dokumente zu ihrem Vortrag oder dem Seminar hochladen, auf die alle anderen Zugriff haben. Weiterhin werden jedem Seminar ein MediaWiki und ein WebDAV-Repository angeboten. Die Erstellung der Seminardokumentation innerhalb der Reihe Dagstuhl Reports wird ebenfalls durch ein Web-Frontend unterstützt, das 2015 komplett neu entwickelt wurde.

Schloss Dagstuhls Internetauftritt[50] bietet nicht nur seinen Gästen, sondern allen Nutzern weltweit Informationen über die folgenden Themen:

- Verbreitung allgemeiner Informationen über das Zentrum, wie Konzept, Programm, Antragsmodalitäten
- Informationen zur Anreise der Teilnehmer, wie Lageplan, Fahrpläne, Taxidienste
- Die Bibliothek mit der Möglichkeit zur Recherche im Dagstuhl-Bibliothekskatalog
- Informationen zu Seminaren und Veranstaltungen, wie Seminarziele, angemeldete Wissenschaftler und Publikationen
- Angebot einer Plattform zum Austausch von Material unter den Seminarteilnehmern

Der Webserver verwaltet die Inhalte mit dem Content-Management-System Typo3. Außer statischen Seiten – nahezu alle in deutschen und in englischen Versionen – werden auch dynamische Seiten angeboten, die über eigene Software generiert werden. So gibt es zu jedem Seminar eine dynamisch generierte Seite, die zu Motivationstext, Teilnehmerliste, Publikationen, etc. führt.

Schloss Dagstuhl offers an increasing number of web-based services to seminar organizers and participants. During the preparation phase, the seminar organizers can check on a daily basis how invited participants are responding to the invitation and which of them have committed to attending. They can also upload a (preliminary) schedule to the seminar web page. All participants can upload seminar- or presentation-related documents to the page, which are then accessible to everyone else. A MediaWiki and WebDAV-related repository are also offered. The making of the seminar documentation inside our Dagstuhl Reports periodical is also supported by a Web-based service that was completely redeveloped in 2015.

In keeping with the center's philosophy, its Internet[50] offerings are not only available to the guests at Dagstuhl but to netizens throughout the world. Objectives and content:

- Dissemination of general information on the center, e.g. concept, program, particulars pertaining to proposal submission
- Offering participants travel information on how to get to the center (site plan, train and bus schedules, taxi services, etc.)
- Presenting the Dagstuhl Informatics Research Library along with its offerings and resources and enabling research in the Dagstuhl library catalogue
- Providing information about seminars and events (e.g. seminar objectives, scientists from whom proposals have been accepted, publications)
- Providing a platform for exchanging materials among seminar participants

The web server administers the content using the Typo3 content management system. Apart from static pages, almost all of which are in German and English, dynamic pages are also offered which are generated by the center's

proprietary software. Each seminar has a dynamically generated page of its own featuring links to a seminar description, list of participants, publications, etc.

9.4 Dagstuhls Küche

Die Mahlzeiten sind ein wichtiger Bestandteil des wissenschaftlichen Programms von Schloss Dagstuhl. Die Sitzordnung wird absichtlich stets zufällig gemischt, um eingefahrene Gruppen aufzuteilen und Gäste zu ermuntern, während ihres Aufenthalts möglichst viele verschiedene Kollegen kennenzulernen. Große Tische im Speiseraum fördern die gemeinschaftliche Interaktion bei den Mahlzeiten.

Dagstuhls Philosophie des Kochens ist einfach: saisonal, gesund und schmackhaft. Unsere Gerichte werden jeden Tag von unseren 10 Mitarbeitern der Küche und unseren Auszubildenden frisch zubereitet. Der Schwerpunkt liegt dabei auf leichtem Essen während des Tages, um unsere Gäste nicht zu ermüden, und auf warmen Gerichten am Abend. Dies steht ein wenig im Widerspruch zur deutschen Tradition, kommt aber der Mehrheit der internationalen Gäste des Zentrums durchaus entgegen.

Sowohl die Zutaten als auch die Gerichte wechseln entsprechend der Saison. An warmen Sommerabenden wird häufig auf der Terrasse vor dem Speisesaal gegrillt, unter anderem saarländische Schwenker, eine lokale Variante des Grillsteaks, die unter dauerndem Schwenken des Grillrostes zubereitet wird. In den kalten Monaten steht einmal wöchentlich ein schmackhafter Eintopf auf dem Speiseplan. Über das Jahr hinweg wird eine ausgewogene Mischung an regionalen und internationalen Spezialitäten aus neuen sowie bewährten und beliebten Rezepten angeboten. Im Allgemeinen sind die angebotenen Gerichte im Sommer etwas leichter und im Winter ein wenig schwerer. Die Küche arbeitet nach dem HACCP-Konzept (Hazard Analysis and Critical Points Concept) und hält sich an die Kennzeichnungspflicht von Allergenen, zu der alle lebensmittelverarbeitenden Betriebe verpflichtet sind. Des Weiteren achten wir auf deklarationsfreie Zusatz- und Konservierungsstoffe.

Alle Gäste, die aus medizinischen oder ethischen Gründen Einschränkungen bei der Speiseauswahl haben, können sich vor dem Seminar bei Schloss Dagstuhl melden. Unsere Küchenmitarbeiter erarbeiten gerne individuelle Lösungen für jeden Gast, soweit es irgendwie möglich ist. Gäste, die koscheres Essen benötigen, haben die Möglichkeit, mitgebrachte abgepackte Speisen selbst zu erhitzen.

Um unseren Gästen trotz eines begrenzten Budgets eine ausgewogene Qualität anbieten zu können, bietet unsere Küche ein Frühstücksbüffet, dienstags bis donnerstags ein Mittagsbuffet sowie ein Menü am Abend an. Montags und freitags wird aus logistischen Gründen auch am Mittag ein Menü serviert. Unser Restaurant mit den großen Fenstern zum Garten des Hauptgebäudes bietet ca. 80 Personen

9.4 *Dagstuhl's Kitchen*

The dining experience at Dagstuhl is an important part of the center's scientific program. Seating arrangements are deliberately mixed 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 at breakfast and lunch.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy, and tasty meals. Everything is freshly prepared each day by the kitchen's 10-person staff and apprentices in training. The focus is on lighter fare during the day in order to aid scientists' concentration, and on a warm meal in the evening, which breaks with the German tradition of a cold evening meal but fits well with the internationality of the center's guests.

Both ingredients and dishes vary with the changing seasons. On warm summer evenings, guests are frequently invited to partake of grilled *Schwenker* (the local variant of barbecued steak) on the outdoor patio adjacent to the dining hall. Warm soups appear weekly on the menu during the colder winter months. 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 to which all food processing establishments are required. Food additives and conservatives for which labeling is non-mandatory are also carefully monitored.

All guests who have due to ethical or health reasons special food requirements can announce their needs previous to the events. Our kitchen staff will then work out individual solutions if at all possible. Guests who need kosher meals can heat up ready-to-eat meals for themselves.

To accomplish all of this within a reasonable budget, the center offers a buffet-style breakfast and a set evening meal served by the kitchen's friendly and engaged staff. From Tuesday to Thursday the kitchen offers a buffet-style lunch. Due to logistical reasons, a set meal is served at lunch on Mondays and Fridays. The large dining-hall, seating up to 80 persons, opens onto the castle garden and patio, and offers a relaxed, familiar atmosphere.

Small and larger breaks punctuate the daily routine. During the small coffee break during the morning hot drinks are served in front of the lecture halls. During the larger one in the afternoon, hot drinks together with freshly baked cake are served in the dinning hall. In addition, there

[50] http://www.dagstuhl.de/

Platz. Hier herrscht eine entspannte und fast familiäre Atmosphäre, was nicht zuletzt auf unsere freundlichen und engagierten Mitarbeiter zurückzuführen ist.

Kleine und große Pausen unterbrechen auf angenehme Weise die tägliche Routine und anstrengenden Diskussionen. In der kleinen Kaffeepause am Vormittag stehen vor den Vortragsräumen heiße Getränke auf einem Kaffeewagen bereit. In der großen Kaffeepause am Nachmittag wird den Gästen im Speiseraum neben heißen Getränken auch frisch gebackener Kuchen angeboten. Darüber hinaus gibt es im Gästehaus, der „alten" Cafeteria und dem Weinkeller jeweils einen Kaffeevollautomaten zur Zubereitung von Kaffee, Kakao und Tee. In der Cafeteria und dem Weinkeller können Gäste Snacks erwerben. Abends gibt es in diesen beiden beliebten Räumen Brot und eine Käseauswahl.

are self-service bean-to-cup coffee machines in the guest house, at the "old" café, and in the wine cellar. Guests can buy small snacks at in the cafe and the wine cellar – two popular after-hours hangouts. Bread and cheese is served there every night.

Kinderbetreuung 9.5 *Childcare*

Schloss Dagstuhl bietet Teilnehmern, die mit Kindern anreisen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser Service kann gegen ein geringes Entgelt im Voraus gebucht werden. Alternativ ist es Eltern auch möglich, eine Begleitperson zur Betreuung des Kindes oder der Kinder mitzubringen. Für Seminarteilnehmer übernimmt Schloss Dagstuhl die Kosten für Verpflegung und Unterkunft der Kinder und der Begleitperson.

Im Jahre 2015 wurden 8 Kinder durch eine Tagesmutter und 17 weitere durch Verwandte betreut. Insgesamt beherbergte Schloss Dagstuhl 25 Kinder von Teilnehmern an 17 Veranstaltungen während 12 Wochen.

Schloss Dagstuhl gladly offers to organize child care with a certified nanny for participants who need to visit our center with young children. The service, which supports families and particularly women computer scientists, can be booked for a small recompense in advance of the seminar. Parents also have the option to bring along their own "nanny," usually a spouse or relative. For seminar participants the costs for room and board are absorbed by the center both both for the children and the accompanying person.

In 2015, Dagstuhl hosted 25 children, 8 of whom were cared for by a nanny on site and 17 by relatives. Participants of 17 events in 12 weeks were thus able to attend although they were travelling with their children.

Freizeit und Ambiente 9.6 *Leisure Facilities*

Die Freizeitanlangen auf Schloss Dagstuhl wurden so gestaltet, dass sie auf unterschiedliche Art und Weise sowohl tagsüber als auch abends die Kommunikation zwischen den Seminarteilnehmern fördern. Die Mischung aus Arbeit und Freizeit in entspannter, familiärer Atmosphäre ist ein wichtiger Bestandteil des Dagstuhl-Konzepts. Gäste leben und arbeiten zusammen in einem Komplex aus drei Gebäuden, im Zentrum das historische Schloss, wo sie rund um die Uhr freien Zugang zu den zahlreichen Freizeiträumen und -anlagen haben. Musikalische Gäste können ihre Fertigkeiten im barocken Musiksaal zu Gehör bringen, wo ein Flügel und diverse andere Instrumente wie z. B. zwei Konzertgitarren zur Verfügung stehen. Unser Zentrum verfügt außerdem über eine Sauna, einen Billardtisch, Tischfußball, Mountainbikes, eine Dartscheibe, einen Freizeitraum mit Fitnessgeräten und Tischtennis sowie einen Außenbereich mit Volleyballnetz.

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, family-style setting is an important part of the Dagstuhl concept. Guests live and work together in a three-building nucleus centered on the historical manor house ("Schloss") 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 Schloss, which features a grand piano and various other instruments, e.g. two concert guitars. Schloss Dagstuhl also has a full sauna, a pool table, table football facilities, mountain bikes, a dartboard, and a recreation room with gym equipment and table tennis as well as outdoor sports grounds featuring a volleyball net.

10 Bibliothek
Research Library

Bestand und Angebot 10.1 *Inventory and Offering*

Die Forschungsbibliothek bildet eines der wichtigsten Angebote. Sie hat sich Dank der Startfinanzierung der Volkswagen-Stiftung und durch zahlreiche Buchspenden von Verlagen und Seminarteilnehmern zu einer der bedeutendsten Informatik-Forschungsbibliotheken in Deutschland entwickelt.

Die Bibliothek erwirbt aktuelle Informatik-Forschungsliteratur thematisch zu den jeweiligen Seminaren, überwiegend in englischer Sprache. Am 31. Dezember 2015 umfasste der Bibliotheksbestand 60 944 bibliographische Einheiten, die vollständig im Online-Katalog verzeichnet sind. Der umfangreiche Zeitschriftenbestand wird fast komplett elektronisch bezogen. Über die DFG-geförderten National- und Allianzlizenzen ermöglicht die Bibliothek Zugriff auf mehrere Tausend weitere elektronische Zeitschriftentitel und Zeitschriftenarchive.

Die Literatur wird in einem attraktiven Bibliotheksturm auf vier Ebenen präsentiert, der auch zahlreiche Leseplätze zum Studium anbietet. Als Präsenzbibliothek steht sie den Dagstuhl-Seminarteilnehmern für ihre Forschungsarbeit vor Ort rund um die Uhr offen. Aber auch externe Wissenschaftler können die Bibliothek nach Voranmeldung nutzen.

Speziell für die Dagstuhl-Seminare werden jede Woche umfangreiche Buchausstellungen präsentiert. Dazu werden alle im Bibliotheksbestand vorhandenen Bücher der Dagstuhl-Seminarteilnehmer für das jeweilige Dagstuhl-Seminar zusammengestellt. An die Autoren ergeht gleichzeitig die Bitte, ihre Bücher zu signieren. Weiter werden alle Buchspenden von Verlagen separat ausgestellt und regelmäßig aktualisiert. Dieser Service wird von Gästen und Verlagen sehr geschätzt.

Überregional steht der Zeitschriftenbestand durch die Teilnahme an der Online-Fernleihe im Rahmen des internationalen Leihverkehrs Bibliotheken aus der ganzen Welt zur Verfügung. Dazu ist der komplette Zeitschriftenbestand sowohl in der Zeitschriftendatenbank (ZDB) als auch in der Elektronischen Zeitschriftenbibliothek (EZB) nachgewiesen.

Auch der Buchbestand ist überregional im Katalog des Südwestdeutschen Bibliotheksverbundes katalogisiert und dadurch international sichtbar.

Seit 2014 ist der Monographienbestand zusätzlich im regional wichtigen „Saarländischen Virtuellen Katalog“ der Saarländischen Universitäts- und Landesbibliothek sichtbar.

Monatlich werden die Metadaten aller Neuerwerbungen der Bibliothek an die Datenbank dblp geliefert.

Die Bibliothek ist Mitglied des regionalen Service LITexpress, einem Lieferdienst rückgabepflichtiger Medien für Bürgerinnen und Bürger in Rheinland-Pfalz, dem Saarland und der deutschsprachigen Gemeinschaft Belgiens. Dabei sollen vor allem die Archivtitel zur Ausleihe bereitgestellt werden.

Über die Internetseite der Bibliothek[51] sind u.a. der Online-Bibliothekskatalog, die Zeitschriftenbestandsliste mit Zugang zu den abonnierten online verfügbaren Zeit-

The Dagstuhl Informatics Research Library is one of the center's most impressive offerings. Thanks to the startup financing by the Volkswagen Foundation and numerous book donations of publishing houses and seminar participants, it numbers among Germany's key informatics research libraries.

The library collects current research literature on informatics topics for the respective seminars, primarily in English. As of December 31, 2015, the library's collection totaled 60,944 bibliographic units, all of which are contained in the online catalog. Almost all scientific journals are provided online only. Apart from subscribed journals, the library also provides access to several thousand other electronic journals and journal archives via the DFG-funded national and alliance licenses.

The literature is arranged on four levels in an attractive library tower, which also offers a large number of recesses for quiet study and research. Being a reference library, it is at the disposal of the Dagstuhl Seminar participants 24/7 for their research work on site. External scholars can also use the library provided they register beforehand.

Especially for all Dagstuhl-Seminars the library weekly arranges comprehensive book exhibits.

All books authored by the participants in the current Dagstuhl Seminars which are available in the library are compiled and presented for each seminar. The authors are kindly asked to sign their books. In addition, all book donations received from publishers are exhibited separately and the exhibits are regularly updated. This service is highly appreciated by the center's guests and publishers alike.

In order to support informatics research in Germany and throughout the world, the center's entire holdings of periodicals are also made available to other libraries, particularly by way of inter-library loan. The library's entire holdings of journals and periodicals are additionally listed in the Zeitschriftdatenbank (ZDB), the world's largest specialized database for serial titles, and in the Electronic Journals Library (EZB).

Nationwide, the library's book holdings are additionally listed in the catalogue of the "Südwestdeutscher Bibliotheksverbund" which allows to search Dagstuhl's library holdings through online catalogs worldwide.

Since 2014 the monographs are also referenced in the regional important "Saarländischer Virtueller Katalog" maintained by the Saarland University and State Library.

Each month the bibliographic metadata of all new library books are delivered to dblp computer science bibliography.

The library is a member of LITexpress, the Virtual Library of Rhineland-Palatinate, Saarland and the German-speaking community of Belgium, a media loan service for the citizens of these three areas. The library's archive items in particular are designed to be made available for loan.

The online catalogue and a comprehensive journal list with access to the subscribed journals as well as other

schriften sowie weitere Informationsangebote der Bibliothek zu erreichen.

information offerings can be accessed via the library's webpage.[51]

Spenden an die Bibliothek 10.2

Die Bibliothek von Schloss Dagstuhl profitiert durch zahlreiche Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2015 Buchspenden von den Verlagen, die in Fig. 10.1 aufgeführt sind. Auch viele Seminarteilnehmer spenden der Bibliothek ihre Bücher. Autorenexemplare, insbesondere von wichtigen, bereits vergriffenen Büchern, werden ebenso dankbar entgegengenommen. Insgesamt erhielt das Zentrum im Berichtszeitraum 795 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. During 2015 the Informatics Research Library received book donations from the publishers listed in Fig. 10.1. The center is also grateful for donations of author's copies, particularly those of major works that are out of print. The center received a total of 795 volumes during the year 2015 as donations from publishing houses and seminar participants.

[51] http://www.dagstuhl.de/de/library/

Birkhäuser Verlag http://www.birkhaeuser-science.com
Eurographics – European Association for Computer Graphics https://www.eg.org
O'Reilly http://www.oreilly.de
Pearson http://www.pearson.de
SIAM – Society for Industrial and Applied Mathematics http://www.siam.org
Springer-Verlag GmbH | Springer Science+Business Media http://www.springer.com

Fig. 10.1
Donations from publishers to the Dagstuhl library.

11 Kunst
Art

Dagstuhl als Galerie 11.1 *Dagstuhl as Art Gallery*

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des schmalen 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 Kontrat 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 Kurator der Kunstsammlung von Schloss Dagstuhl. Das Zentrum veranstaltet jährlich etwa drei bis vier Kunstausstellungen für jeweils zwei bis drei Monate. Die zwei Ausstellungen (siehe Fig. 11.1), die im Jahr 2015 stattfanden, sind nachfolgend beschrieben. Die jeweils aktuellen Ausstellungen sind nach Anmeldung auch für die interessierte Öffentlichkeit zugänglich.

Vera Loos, »Something is happening here | But I don't know what it is«

Die Ausstellung mit den Werken der in Saarbrücken lebenden Künstlerin Vera Loos fand vom 20. April bis 24. Juli 2015 auf Schloss Dagstuhl statt. In Saarlouis geboren, studierte Vera Loos zunächst in Saarbrücken Sprachwissenschaft, Kunstgeschichte und Germanistik. Beginnend ab 1980 machte sie dann eine künstlerische Ausbildung, unter anderem bei Marina Hartwahn, Natascha Pop, Tina Stein und Wolf Werdigier. Ausstellungen von Frau Loos fanden nicht nur deutschlandweit, sondern auch in Frankreich und China statt. Eines ihrer ausgestellten Bilder ist in Fig. 11.2 auf Seite 216 gezeigt.

Vera Loos schreibt über ihre Bilder: *Das Absurde in der menschlichen Existenz beschäftigt mich und ist mir immer wieder eine Quelle der Inspiration. Die skurrilen Konstellationen des Alltags, die Sinnsuche im Unsinnigen, das Gefühl der Verlorenheit sind Themen, die ich bildlich darzustellen versuche. Ich wünsche mir, dass meine Bilder im Betrachter einen Strom von Spekulationen auslösen, die zu Fantasien über das Bildgeschehen und die Befindlichkeit meiner Protagonisten führen. Meine Titel sind häufig Zitate aus der Literatur, die mir Bildideen liefert und mich beflügelt.*

Exhibitions of artists 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 walls of the narrow gallery, the artworks offset the otherwise ascetic nature of the new building. These temporary exhibits offer a fresh and dynamic counterpoint to center's permanent collection, which can be found scattered throughout the three buildings.

Prof. Reinhard Wilhelm continued to serve as curator of the Schloss Dagstuhl art collection following his retirement as Scientific Director of the center in April 2014. The center holds approximately three to four art exhibits per year, with each exhibit generally running for two to three months. The two exhibits (cf. Fig. 11.1) hosted by Schloss Dagstuhl in 2015 are described below. Current exhibitions are open to the interested public upon request.

Vera Loos, »Something is happening here | But I don't know what it is«

Featuring the works of Saarbrücken-based artist Vera Loos, this exhibition was on display at Schloss Dagstuhl from April 20 to July 24, 2015. Loos, who was born in Saarlouis, Germany, originally studied linguistics, art history, and German philology. Starting in 1980, she trained with Marina Hartwahn, Natascha Pop, Tina Stein, and Wolf Werdigier, amongst others. Her works have been displayed not only throughout Germany, but also in France and China. One of her exhibited paintings is shown in Fig. 11.2 on Page 216.

Vera Loos writes the following about her pictures: *What matters to me and continues inspire me is the absurdity of human existence. The bizarre constellations of day-to-day life, the search for meaning in meaninglessness, the sensation of forlornness are themes I strive to depict. I wish for my pictures to trigger a stream of speculations leading to images about the image narrative and the sensitivities/inner state of my protagonists. My title soften are quotations from the literature that provides me with ideas for pictures and inspires me.*

Vera Loos, »Something is happening here But I don't know what it is«
Works by artist Vera Loos | April 20 to July 24, 2015

Gudrun Emmert, »colored promenade«
Works by artist Gudrun Emmert | September 14 to December 18, 2015

Fig. 11.1
Art exhibitions in 2015.

Gudrun Emmert, »colored promenade«

Die Ausstellung mit den Werken der in Saarbrücken lebenden Künstlerin Gudrun Emmert fand vom 14. September bis 18. Dezember 2015 auf Schloss Dagstuhl statt. In Münnerstadt (Unterfranken) geboren, studierte Gudrun Emmert bei Professor Lobeck an der damaligen Gesamthochschule Kassel Freie Kunst (Malerei). Ihre Kunstwerke werden seit 1986 in Einzelaustellungen und Ausstellungsbeteiligungen in ganz Deutschland gezeigt.

Die Malerin Gudrun Emmert folgt in ihren Bildern einem formalen Schema, etwa Streifen oder Gewebe. Dennoch gibt es einen gewissen Spielraum, in dem sich die Bilder frei entwickeln können. Ihr eigentliches Interesse gilt dabei dem Malvorgang: Farben, Formen und Flächen überlagern sich, werden geschichtet, Gegenständliches wird nur angedeutet und löst sich teilweise wieder in der autonomen Malerei auf. Dabei bleibt der Entstehungsprozess eines Bildes für den Betrachter nachvollziehbar.

Gudrun Emmert, »colored promenade«

This exhibition, featuring the works of Saarbrücken-based artist Gudrun Emmert, took place at Schloss Dagstuhl from September 14 to December 18, 2015. Born in Münnerstadt, Germany, Gudrun Emmert studied Free Arts (painting) with Professor Lobeck at the former Gesamthochschule Kassel. Her Works have been on display in solo exhibitions as well as exhibition participations throughout Germany.

In her pictures, Gudrun Emmert follows a formal scheme, for instance stripes or weave. Nonetheless, there is a certain leeway so that the pictures may develop free of restriction. The artist's main interest is the act of painting itself: colours, shapes and surfaces overlap or are layered, figures are only implied and partially dissolve within the autonomous painting. Thus, the origination process of a picture remains comprehensible for the viewer.

11.2 Kunstankauf durch Spenden / *Art Sponsorship and Donations*

Das Internetangebot von Schloss Dagstuhl enthält eine Seite, die es Teilnehmern, Einzelpersonen und Gruppen ermöglicht, Kunst für Dagstuhl zu stiften. Die Kunstobjekte werden über das Internet angeboten, dabei wird der Preis in kostengünstige Anteile aufgeteilt. Sobald alle Anteile eines Bilds gezeichnet sind, werden die Teilnehmer aufgefordert, den Gegenwert der bestellten Anteile als Spende einzuzahlen, wodurch dann das Objekt angekauft werden kann. Die Stifter werden sowohl in der virtuellen Internet-Galerie von Schloss Dagstuhl als auch an dem realen Objekt genannt. Dadurch ist es Schloss Dagstuhl möglich, Werke von Künstlern, die im Zentrum ausgestellt haben, anzukaufen und auszustellen. Auf diesem Weg konnte 2015 die Finanzierung zweier Kunstwerke abgeschlossen werden und die Bilder zur permamenten Ausstellung in Schloss Dagstuhl hinzugefügt werden.

Darüber hinaus spendete Helen C. Purchase, langjährige Teilnehmerin verschiedener Dagstuhl-Seminare, dem Zentrum ein Bild (siehe Fig. 11.3 auf Seite 217) aus der Ausstellung »colored promenade« von Gudrun Emmert. Dr. Purchase bedankt sich damit auch für den positiven Einfluss der Dagstuhl-Seminare auf ihre Karriere. Im Jahr 2015 erhielt Schloss Dagstuhl insgesamt 745 Euro von verschiedenen Spendern.

Nähere Informationen und aktuelle Neuigkeiten finden sich auf der Kunst-Webseite[52] von Dagstuhl. Wir möchten diese Stelle nutzen, allen Spendern, die 2015 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

Dagstuhl's website contains a page featuring an Internet gallery enabling participants, individuals, and groups to make contributions to Dagstuhl for art donations. The works of art are featured on the Internet and donations are made by acquiring shares at affordable prices. Donors pay the value of their pledged shares as soon as a piece is fully subscribed for, thus allowing it to be purchased. Donors' names appear in Dagstuhl's online art gallery and also on the art items themselves. The art donation program also benefits the center, enabling Schloss Dagstuhl to purchase works of art from those who exhibit at the center. In 2015 Dagstuhl could finish the financing of two images. These images are now part of Schloss Dagstuhl's permanent art display.

Additionally, Helen C. Purchase, long-standing participant of several Dagstuhl Seminars, donated one of the pictures (see Fig. 11.3 on Page 217) from Gudrun Emmert's exhibition »colored promenade« von Gudrun Emmert. Dr. Purchase also wanted to express her thanks for Schloss Dagstuhl's positive influence on her career. In 2015, Schloss Dagstuhl received a total of 745 Euro from several donors for the center's art collection.

For further information about the art exhibits at Schloss Dagstuhl and the art program in general, please visit Dagstuhl's art webpage[52]. We would like to take this opportunity to thank all of those who generously donated to the center's art collection in 2015.

[52] http://www.dagstuhl.de/art/

Dagstuhls Permanente Kunstausstellung

11.3

Dagstuhl Permanent Art Exhibition

Die von Gästen immer wieder positiv hervorgehobene Kunstsammlung geht auf den Gründungsdirektor Professor Wilhelm zurück. Seine Idee war es, den 1995 neueröffneten Speisesaal und den etwa ein Jahr älteren Neubau, durch Kunstwerke zu beleben. Dazu startete er die oben beschrieben Kunstaustellungen. Unter Mitwirkung der Künstler wird aus jeder Ausstellung ein Werk ausgewählt, für das dann Spender gesucht werden. In den letzten 20 Jahren kamen so ungefähr 180 Kunstwerke zusammen. Auch durch diese Initiative angeregt und verstärkt erhielt Dagstuhl in den vergangenen Jahren weitere Spenden von Künstlern und Mäzenen. Die Arbeiten kommen in den Räumen des Zentrums in Wadern sowie in der Geschäftsstelle in Saarbrücken sehr gut zur Geltung.

The art collection, continually praised by guests, was intiated by Founding Director Professor Wilhelm. It was his idea to use works of art in order to enliven the New Building as well as the dining room opened in 1994 and 1995, respectively. To this end, Professor Wilhelm launched the exhibitions described above. Assisted by the artists, one picture from each exhibition was chosen and donors were drummed up. Thus, approximately 120 works of art could be acquired over the last 20 years. Additionally, this initiative has increasingly encouraged artists and patrons to make donations. All of the pictures adorn the rooms of Schloss Dagstuhl in Wadern as well as the Dagstuhl Office in Saarbrücken.

Fig. 11.2
"Mit Links" painting by Vera Loos shwon in her exhibit at Schloss Dagstuhl.

Fig. 11.3
"Nr. 7" painting by Gudrun Emmert, donated by Helen C. Purchase to Schloss Dagstuhl.

12 Struktur der Gesellschaft
Structure of the Company

Gründung und Gesellschafter 12.1 *Formation and Shareholders*

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit derzeit elf Gesellschaftern (siehe Fig. 12.1) organisiert. Dies sind die vier Gesellschafter, die Schloss Dagstuhl gegründet haben, nämlich die Gesellschaft für Informatik e. V. (GI), die Universität des Saarlandes, die Technische Universität Kaiserslautern und das Karlsruher Institut für Technologie (KIT). Als vier weitere Gesellschafter wurden 1994 die Technische Universität Darmstadt, die Johann Wolfgang Goethe-Universität Frankfurt am Main, die Universität Stuttgart und die Universität Trier aufgenommen. Drei international renommierte Forschungsinstitute, das Institut National de Recherche en Informatique et en Automatique (INRIA, Frankreich), das Centrum Wiskunde & Informatica (CWI, Niederlande) und die Max-Planck-Gesellschaft (MPG, Deutschland) wurden 2005/2006 als weitere Gesellschafter aufgenommen.

Aufgrund eines Beschlusses der Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung (heute Gemeinsame Wissenschaftskonferenz) wurde das Zentrum mit Wirkung zum 1. Januar 2006 als Serviceeinrichtung für die Forschung in die gemeinsame Forschungsförderung von Bund und Ländern aufgenommen. Es ist seit 2005 Mitglied der Leibniz-Gemeinschaft. Entsprechend wurde 2008 der Name des Zentrums von vormals „Internationales Begegnungs- und Forschungszentrum für Informatik“ in „Schloss Dagstuhl – Leibniz-Zentrum für Informatik“ geändert.

Schloss Dagstuhl wurde im Juli 2009 erstmals durch die Leibniz-Gemeinschaft evaluiert. Die Stellungnahme der Evaluierungs-Kommission vom März 2010 war sehr positiv: Schloss Dagstuhl widme sich mit herausragendem Erfolg seiner Aufgabe, die internationale Informatikforschung mit einem Seminarzentrum für wissenschaftliche Veranstaltungen zu unterstützen.

Schloss Dagstuhl is operated as a non-profit organization by currently eleven associates (cf. Fig. 12.1), including its four founding associates: the Gesellschaft für Informatik e. V.[53] (GI), the Universität des Saarlandes, the Technische Universität Kaiserslautern, and the Karlsruher Institut für Technologie (KIT). In 1994, the organization was extended to include four new associates: the Technische Universität Darmstadt, the Johann Wolfgang Goethe-Universität Frankfurt am Main, the Universität Stuttgart and the Universität Trier. Finally, in 2005 and 2006, three internationally renowned research institutes joined the association: the Institut National de Recherche en Informatique et en Automatique (INRIA, France), the Centrum Wiskunde & Informatica (CWI, Netherlands), and the Max-Planck-Gesellschaft (MPG, Germany).

By resolution of the Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung[54] (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”[55] to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”[56] 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.

Organe der Gesellschaft 12.2 *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.

Die Gesellschafterversammlung

Die Gesellschafter beschließen über alle Änderungen an der Gesellschaft, insbesondere über die Aufnahme

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.

Shareholders’ Meeting

All changes to the company, in particular the inclusion of new associates, the revision of the Shareholders’

[53] engl.: German Informatics Society
[54] engl.: Federal/State Government Commission for Educational Planning and Research Promotion
[55] engl.: International Conference and Research Center for Computer Science
[56] engl.: Schloss Dagstuhl – Leibniz Center for Informatics

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 einstimmig 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. 12.2) 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 Amtszeit der Aufsichtsratsmitglieder beträgt vier volle abgeschlossene Geschäftsjahre, was de facto einer fünfjährigen Amtszeit entspricht. Die Vertreter der Universitäten in Darmstadt und Stuttgart wechseln 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, die gemeinsam die Gesellschaft vertreten. Die Geschäftsführung besteht aus dem *Wissenschaftlichen Direktor* und dem *Technisch-administrativen Geschäftsführer*.

Der Wissenschaftliche Direktor ist verantwortlich für die wissenschaftlich-fachliche Zielsetzung und die Programmgestaltung. Er ist Mitglied des Wissenschaftlichen Direktoriums und leitet dieses. Seit Mai 2014 ist Prof. Raimund Seidel, Ph. D., der wissenschaftliche Direktor von Schloss Dagstuhl.

Der Wissenschaftliche Direktor wird dem Aufsichtsrat von einer Findungskommission zur Berufung vorgeschlagen. Dieser Findungskommission gehören mindestens der Vorsitzende des Aufsichtsrats und der Vorsitzende des

agreement and the dissolution of the company, are decided by the shareholders. Shareholders also confirm new members forwarded by them to the Supervisory Board and the appointment or recall of the managing directors. In accordance with their shares, all shareholders currently have the same number of votes except the Gesellschaft für Informatik, which has three times the number of votes of other shareholders in proportion to its larger number of shares. Decisions are made in shareholders' meetings which take place at least once the year, or via a written vote with unanimous consent.

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. 12.2) is composed of four representatives of the Gesellschaft für Informatik, one representative each of the three founding universities, two representatives of the four universities that subsequently joined, and one representative each of the German federal government and the two host state governments of Saarland and Rhineland-Palatinate. The members of the Supervisory Board hold office for four full fiscal years, which are, de facto, five years. 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 Advisory 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 who jointly represent the company. These are the *Scientific Director* and the *Technical Administrative Director*.

The Scientific Director is in charge of drafting the company's scientific goals and program planning. He is a member and the chairperson of the Scientific Directorate. Since May 2014, Prof. Raimund Seidel, Ph.D., is the Scientific Director of Schloss Dagstuhl.

The Supervisory Board appoints the Scientific Director on basis of the recommendation of a selection committee 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.

Wissenschaftlichen Beirats an. Die Amtszeit des Wissenschaftlichen Direktors beträgt fünf Jahre.

Die technischen und administrativen Aufgaben werden vom Technisch-administrativen Geschäftsführer wahrgenommen. Seit Juli 2014 hat Frau Heike Meißner diese Position inne. Siehe Fig. 12.3.

Gremien der Gesellschaft 12.3

Die Organe von Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH werden durch drei Gremien unterstützt. Es sind die folgenden:

- Wissenschaftliches Direktorium
- Wissenschaftlicher Beirat
- Kuratorium

Details zu den Gremien werden in den folgenden Abschnitten ausgeführt.

Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 12.4) ist für die Realisierung des Gesellschaftszwecks in fachlich-wissenschaftlicher Hinsicht verantwortlich. Es hat das Forschungs- und Veranstaltungsprogramm der Gesellschaft festzulegen, seine fachlich-wissenschaftliche Qualität zu sichern und seine Durchführung zu überwachen. Als wesentlicher Bestandteil dieser Aufgabe werden die Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von Mitgliedern des Wissenschaftlichen Direktoriums begutachtet. Auf den zweimal im Jahr stattfindenden Direktoriumssitzungen werden die Anträge diskutiert und es wird über ihre Annahme entschieden.

Der Wissenschaftliche Direktor gehört dem Wissenschaftlichen Direktorium an. Er empfiehlt dem Aufsichtsrat die Größe des Direktoriums. Neben den Gesellschaftern können das bestehende Wissenschaftliche Direktorium sowie der Beirat Kandidaten für das Wissenschaftliche Direktorium benennen. Die Auswahl der Kandidaten, die dem Aufsichtsrat zur Ernennung vorgeschlagen werden, obliegt dem Beirat zusammen mit dem Wissenschaftlichen Direktor.

Die Amtszeit der Mitglieder des Wissenschaftlichen Direktoriums – mit Ausnahme der des Wissenschaftlichen Direktors – beträgt drei Jahre. Sie beginnt am 1. November des Jahres ihrer Berufung und endet drei Jahre später am 31. Oktober. Eine Wiederberufung ist auch mehrfach möglich.

Der Wissenschaftliche Beirat

Die Aufgaben des Wissenschaftlichen Beirats 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 Forschungs- und Entwicklungsplanung, nimmt Stellung zu den Programmbudgets und gibt Empfehlungen zum Ressourceneinsatz. Er unterstützt weiterhin den Aufsichtsrat bei wichtigen Entschei-

The Technical Administrative Director is responsible for technical and administrative tasks. Since July 2014, Ms Heike Meißner holds this position. See Fig. 12.3.

Dagstuhl Bodies 12.3

The organs of Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH are supported by the following bodies:

- Scientific Directorate
- Scientific Advisory Board
- Industrial Advisory Board

Detailed information about these boards can be found in the sections below.

Scientific Directorate

The Scientific Directorate (see Fig. 12.4) is responsible for carrying out the company objectives from a technical and scientific point of view. It must determine the research and event program, ensure its technical and scientific quality, and monitor its execution. As a main task in support of this objective, members of the Scientific Directorate review proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In its twice-yearly directorate meetings, the Scientific Directorate discusses the proposals and decides which of them to accept or reject.

The Scientific Director is member of the Scientific Directorate. He recommends to the Supervisory Board the number of Scientific Directorate members. Candidates for the Scientific Directorate may be suggested not only by the shareholders, but also by the Scientific Directorate and the Scientific Advisory Board. The selection of candidates, which are recommended to the Supervisory Board for appointment, is carried out by the Scientific Advisory Board together with the Scientific Director.

The term of office of Scientific Directorate members – with the exception of the Scientific Director – is three years. It begins on November 1 of the year of appointment and ends three years later on October 31. Multiple reelections are possible.

Scientific Advisory Board

The tasks of the Scientific Advisory Board 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 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

dungen zur Weiterentwicklung von Schloss Dagstuhl und bei der Gewinnung von Leitungspersonal. Zum anderen führt der Wissenschaftliche Beirat mindestens einmal zwischen je zwei Evaluierungen durch den Senatsausschuss Evaluierung der Leibniz-Gemeinschaft ein Audit durch, bei dem die gesamte Einrichtung begutachtet wird. Ein Bericht über das Audit wird der Leitung, dem Aufsichtsrat und dem Senatsausschuss vorgelegt.

Der Wissenschaftliche Beirat (siehe Fig. 12.5) 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. 12.6) 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.

Supervisory Board in important decisions with regard to future development of the institute as well as the acquisition of management staff. On the other hand, it carries out an audit of the entire institute between two evaluations by the Senatsausschuss Evaluierung of the Leibniz Association. A report on this audit is sent to the management, the Supervisory Board and the Senatsausschuss.

The Scientific Advisory Board (see Fig. 12.5) 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 Advisory Board

The Industrial Advisory Board (see Fig. 12.6) 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 Advisory 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 Advisory 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. 12.1
Associates.

Aufsichtsrat | Supervisory Board
Dr. Doreen Becker Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government
Prof. Alejandro P. Buchmann, Ph. D. Technische Universität Darmstadt, Germany | Representative of Technische Universität Darmstadt
Dr. Peter Federer Gesellschaft für Informatik e. V., Bonn, Germany | Representative of Gesellschaft für Informatik e. V. | *tenure ended in June 2015*
Dr. Christian Heimann Ministerium für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of Rhineland-Palatinate state
Prof. Dr.-Ing. Dr. h. c. Stefan Jähnichen Technische Universität Berlin, Germany | Representative of Gesellschaft für Informatik e. V. | Chairman of the Supervisory Board
Prof. Dr.-Ing. Peter Liggesmeyer Technische Universität Kaiserslautern und Fraunhofer IESE, Germany | Representative of Gesellschaft für Informatik e. V.
Prof. Dr. Volker Linneweber Universität des Saarlandes, Saarbrücken, Germany | Representative of Universität des Saarlandes
Prof. Dr. Erhard Plödereder Universität Stuttgart, Germany | Representative of Universität Stuttgart
Prof. Dr. Arnd Poetzsch-Heffter Technische Universität Kaiserslautern, Germany | Representative of Technische Universität Kaiserslautern
Alexander Rabe Gesellschaft für Informatik e. V., Berlin, Germany | Representative of Gesellschaft für Informatik e. V. | *tenure started in June 2015*
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany | Representative of the Saarland
Prof. Dr. Peter H. Schmitt Karlsruher Institut für Technologie, Germany | Representative of Karlsruher Institut für Technologie
Prof. em. Dr.-Ing. Dr.-Ing. h. c. Roland Vollmar Karlsruher Institut für Technologie, Germany | Representative of Gesellschaft für Informatik e. V.

Fig. 12.2
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. 12.3
Management.

Wissenschaftliches Direktorium | Scientific Directorate
Prof. Gilles Barthe, Ph. D. IMDEA Software Institue, Madrid, Spain | *tenure started in November 2015*
Prof. Dr. Bernd Becker Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Stefan Diehl Universität Trier, Germany
Prof. Dr. Hans Hagen Technische Universität Kaiserslautern, Germany
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany
Prof. Dr.-Ing. Oliver Kohlbacher Eberhard Karls Universität Tübingen, Germany
Dr. Stephan Merz Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy – Grand Est, France
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany
Prof. Dr. Bernhard Nebel Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Bernt Schiele Max-Planck-Institut für Informatik, Saarbrücken, Germany
Prof. Dr. Nicole Schweikardt Humboldt-Universität zu Berlin, Germany
Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Saarbrücken, Germany
Prof. Dr. Ir. Arjen P. de Vries Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands
Prof. Dr. Michael Waidner Technische Universität Darmstadt, Germany | *tenure ended in October 2015*
Prof. Dr.-Ing. Klaus Wehrle Rheinisch-Westfälische Technische Hochschule Aachen, Germany
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Universität des Saarlandes, Saarbrücken, Germany

Fig. 12.4
Scientific Directorate.

Wissenschaftlicher Beirat | Scientific Advisory Board
Prof. Dr. Christel Baier Technische Universität Dresden, Germany | *tenure started in June 2015*
Prof. Dr. Manuel V. Hermenegildo IMDEA Software Institute, Madrid and Technical University of Madrid, Spain
Prof. Dr. Claude Kirchner Institut National de Recherche en Informatique et en Automatique (INRIA), Villers-lès-Nancy, France
Prof. Dr. Friedhelm Meyer auf der Heide Heinz Nixdorf Institute, Paderborn and Universität Paderborn, Germany | since March 2015 Chairman of the Scientific Advisory Board
Prof. Dr.-Ing. Dr. h. c. Andreas Reuter HITS GmbH, Heidelberg, Germany
Prof. em. Dr. Dr. h. c. Otto Spaniol RWTH Aachen, Germany
Dr. Susanne Reichrath Staatskanzlei des Saarlandes, Saarbrücken, Germany | *Guest*

Fig. 12.5
Scientific Advisory Board.

Kuratorium | Industrial Advisory Board
Dr. Udo Bub EIT ICT Labs, Berlin, Germany
Dr. Jorge R. Cuéllar Siemens AG, München, Germany | *tenure ended in May 2015*
Dr.-Ing. Elmar Dorner SAP Research, Karlsruhe, Germany
Dr. Jo Ebergen Oracle Labs, Redwood Shores, United States
Dr.-Ing. Uwe Franke Daimler AG, Böblingen, Germany | *tenure started in June 2015*
Dr. Goetz Graefe HP Labs, Palo Alto, United States
Prof. Dr. Ralf Guido Herrtwich Daimler AG, Böblingen, Germany | *tenure ended in May 2015*
Prof. Dr. Ulrich Lauther Siemens AG, München, Germany | *tenure ended in May 2015*
Dr. Michael May Siemens AG, München, Germany | *tenure started in June 2015*
Prof. Dr. Prabhakar Raghavan Google Inc. and Consulting Professor at Stanford University, United States | *tenure ended in May 2015*
Prof. Dr.-Ing. Dr. h. c. Andreas Reuter HITS GmbH, Heidelberg, Germany | *tenure ended in May 2015*
Prof. Dr. Volker Tresp Siemens AG, München, Germany and Ludwig-Maximilians-Universität München, Germany | *tenure ended in May 2015*
Dr. Andreas Wierse SICOS BW GmbH, Stuttgart, Germany | *tenure started in June 2015*

Fig. 12.6
Industrial Advisory Board.

13 Förderverein “Freunde von Dagstuhl”
Association “Friends of Dagstuhl”

Förderverein "Freunde von Dagstuhl"
Association "Friends of Dagstuhl"

Holger Hermanns (Universität des Saarlandes, Germany)
Erich Reindel (Universität des Saarlandes, Germany)

Seit dem 6. Mai 2014 haben die Freunde von Dagstuhl endlich eine Heimat. An diesem Tag haben sich 16 Freunde von Schloss Dagstuhl zusammen gefunden, um den Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V. zu gründen. Der sehr technische und holprig klingende Name spiegelt aber exakt den Vereinszweck wider: die Förderung von Wissenschaft und Forschung im Leibniz-Zentrum für Informatik in Schloss Dagstuhl. Für die Webpräsenz wurde mit „Friends of Dagstuhl" ein wesentlich geschmeidigerer Name gewählt (http://www.friends-of-dagstuhl.de).

Der Verein ist darauf ausgerichtet, finanzielle Mittel zur erfolgreichen Umsetzung des Vereinszweck zu beschaffen und bereitzustellen sowie treuhänderisch die zu diesem Zweck anvertrauten Mittel zu verwalten. Die Stiftung Informatikzentrum Schloss Dagstuhl wurde daher auch als nicht rechtsfähige Stiftung in den Verein überführt. Seit Ende 2014 vertreten nun die Freunde von Dagstuhl die Stiftung im Rechts- und Geschäftsverkehr und verwalten das Stiftungsvermögen. Der Verein wird geleitet von dem Vorstandsvorsitzenden Prof. Holger Hermanns von der Universität des Saarlandes und seiner 1. Stellvertreterin Angelika Müller-von Brochowski (Schriftführerin), die von 1991 bis zum Eintritt in ihren Ruhestand im Jahr 2012 in der Geschäftsstelle von Schloss Dagstuhl tätig war. Komplettiert wird der Vorstand von Erich Reindel (Schatzmeister), dem Geschäftsführer der Fachrichtung Informatik der Universität des Saarlandes. Siehe Fig. 13.1.

Das Jahr 2015 wurde vom Vereinsvorstand genutzt, um die notwendige Infrastruktur aufzubauen und Mitglieder zu werben. Im Juli 2015 konstituierte sich auch der Stiftungsrat, dem die grundsätzlichen Richtungsentscheidungen bei der Verwaltung und Verwendung des Stiftungsvermögens obliegt. Mitglied dieses ersten Stiftungsrates ist satzungsgemäß der Vorstandsvorsitzende des Treuhänders, Prof. Holger Hermanns. Weitere Mitglieder sind Prof. Kurt Mehlhorn vom Max-Planck-Institut für Informatik (MPII), Saarbrücken, sowie Prof.in Dorothea Wagner vom Karlsruher Institut für Technologie (KIT).

Inzwischen gehören dem Verein 35 persönliche sowie 4 Institutionelle Mitglieder an. Gerade im Hinblick auf die noch geringe Anzahl institutioneller Mitglieder wünschen sich die Freunde von Dagstuhl noch regen Zulauf.

As of May 6, 2014, Dagstuhl supporters finally have a home. On that day, 16 friends of Dagstuhl gathered in order to found the registered association in support of Schloss Dagstuhl – Leibniz Center for Informatics (Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e.V.). This very technical and rather clumsy name nevertheless reflects the precise purpose of the association: the support of science and research at the Leibniz Center for Informatics at Schloss Dagstuhl. A significantly smoother name, i.e. "Friends of Dagstuhl", was chosen for the website (http://www.friends-of-dagstuhl.de).

The association aims at acquiring and providing funds for the successful execution of its purpose, as well as holding these funds in trust. The Dagstuhl Foundation (Stiftung Informatikzentrum Schloss Dagstuhl) was therefore integrated into the association as a dependent foundation. Since late 2014, Friends of Dagstuhl represent the foundation in legal and business transactions and manage the foundation assets. The association is chaired by Prof. Holger Hermanns, Saarland University, and the first deputy chairperson, Angelika Müller-von Brochowski (secretary), who worked at the Dagstuhl Office from 1991 until her retirement in 2012. The third board member is Erich Reindel (treasurer), executive of the Computer Science Department at Saarland University. See Fig. 13.1.

In 2015, the association board established the necessary infrastructure and recruited members. A foundation council was was constituted in July of the same year. The council is responsible for general trend-setting decisions in terms of managing and appropriating foundation assets. According to the statutes, chairman Prof. Holger Hermanns is a member of the council. Other members are Prof. Kurt Mehlhorn, Max Planck Institute for Informatics (MPII), Saarbrücken, and Prof. Dorothea Wagner, Karlsruhe Institute of Technology (KIT).

By now, the association has 35 individual and 4 institutional members. Especially with regard to the small number of institutional members, Friends of Dagstuhl look forward to welcoming new members.

Fig. 13.1
Der Vorstand des Vereins “Friends of Dagstuhl”, v.l.n.r.: Prof. Holger Hermanns, Angelika Müller-von Brochowski, und Erich Reindel.
The chair of the association “Friends of Dagstuhl”, f.l.t.r.: Prof. Holger Hermanns, Angelika Müller-von Brochowski, and Erich Reindel.

14 Statistik 2015
Statistics 2015

Statistiken zum wissenschaftlichen Programm

14.1

Statistics on the Scientific Program

In diesem Abschnitt werden statistische Daten zum wissenschaftlichen Programm und der Zusammenstellung der Teilnehmer aufgeführt. Die Diagramme und Tabellen sind dabei wie nachfolgend beschrieben gegliedert.

Antrags-bezogene Daten: Die Anzahl eingereichter Anträge von Dagstuhl Seminaren und Dagstuhl Perspektiven Workshops sowie deren Akzeptanzraten sind in Fig. 14.1 dargestellt. Fig. 14.2 zeigt, wie die akzeptierten Seminare und Workshops sich bezüglich Größe und Länge aufgliedern.

Veranstaltungs-bezogene Daten: Fig. 14.3 zeigt Anzahl und Anteil der eingeladenen Seminarteilnehmer, welche die Einladung annehmen bzw. ablehnen. Die Verteilung dieser Annahmerate ist in Fig. 14.4 dargestellt. Fig. 14.5 zeigt dagegen, wie viel Prozent der zugesagten Größe (gemessen an der Personenanzahl) tatsächlich von einem Seminar belegt wurde. Daten zu Anzahl, Größe und Dauer der durchgeführten Seminare sind in Fig. 14.6 angegeben. Fig. 14.7 zeigt die Anzahl der verschiedenen Veranstaltungstypen.

Teilnehmer-bezogene Daten: Die Teilnehmerzahlen – abhängig vom Veranstaltungstyp – gibt Fig. 14.8 an. Fig. 14.9 zeigt die Verteilung der Herkunftsländer unserer Gäste.

Umfrage-bezogene Daten: Hier stellen wir ausgewählte Daten unserer fortlaufenden Befragung von Teilnehmern an Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops dar. Ein Überblick über die Ergebnisse der regelmäßigen Gästebefragungen kann Fig. 14.10 entnommen werden. Die Anzahl von früheren Seminarbesuchen kann man Fig. 14.11 entnehmen. Fig. 14.12 gibt Auskunft über die Altersstruktur der Teilnehmer. Während Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops sich primär an Forscher aus Universitäten und Forschungseinrichtungen richten, sind auch Anwender und Forscher aus der Industrie stets willkommen. Die Verteilung ihres Anteils ist in Fig. 14.13 gezeigt.

Auslastungs-bezogene Daten: Die Auslastung des Zentrums wird schließlich in Fig. 14.14 an Hand der Gasttage und ihrer Verteilung über die einzelnen Wochen getrennt nach Veranstaltungstypen aufgezeigt.

Geschlechter-bezogene Daten: Fig. 14.15 enthält Daten zur Geschlechter-Verteilung in der Seminarleitung. Dagegen zeigt Fig. 14.16 die Quote von Frauen bei der Beantragung von Seminaren sowohl bezüglich der Teams als auch bezüglich der gesamten Antragsteller. Die Abbildungen Fig. 14.17 und Fig. 14.18 zeigen insbesondere die Anteile weiblicher Teilnehmer bzw. Einladungen an weibliche Wissenschaftler. Die Verteilung der Rate der weiblichen Teilnehmer in den einzelnen Seminaren wird in 14.19 aufgezeigt.

This section provides statistical data about the scientific program and the composition of program participants. Charts and tables in this chapter may be outlined as follows.

Proposal-related data: Fig. 14.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, as well as acceptance rates for recent years. The size and duration of accepted seminars and workshops are displayed in Fig. 14.2.

Event-related data: Fig. 14.3 shows the number and the fraction of invited seminar participants who accepted or declined the invitation. The distribution of the rate is given in Fig. 14.4. In contrast, Fig. 14.5 visualizes how much of the reserved space was actually used by seminar participants. Data related to the number of seminars held in the last years together with their sizes and durations are given in Fig. 14.6. Fig. 14.7 shows the distribution of different types of events at Dagstuhl.

Participant-related data: Fig. 14.8 shows the number of participants according to event type. Fig. 14.9 shows the distribution of country affiliations.

Survey-related data: In this section we present data obtained from our ongoing Dagstuhl Seminar and Dagstuhl Perspectives Workshop guest survey project. An overview of the results of the participant survey for Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in Fig. 14.10. Fig. 14.11 displays how often participants have attended seminars in the past. Fig. 14.12 gives data on the seniority of participants. While Dagstuhl Seminars and Dagstuhl Perspectives Workshops are mainly oriented towards academic researchers also researchers and developers from industry are welcome. The distribution of their fraction compared to all participants of a seminar is shown in Fig. 14.13.

Utilization-related data: Finally, Fig. 14.14 states the number of guest days – separated by event type – hosted at Schloss Dagstuhl as well as their distribution about the weeks.

Gender-related data: Fig. 14.15 shows mixed-gender data with respective to organizer teams of Dagstuhl Seminars and Dagstuhl Perspectives Workshops. In contrast Fig. 14.16 presents these data with respect to proposed seminar events. In Fig. 14.17 and Fig. 14.18 data is given with regard to female participants and invitees, respectively. The distribution of the rate of female participants by seminar and year is displayed in Fig. 14.19.

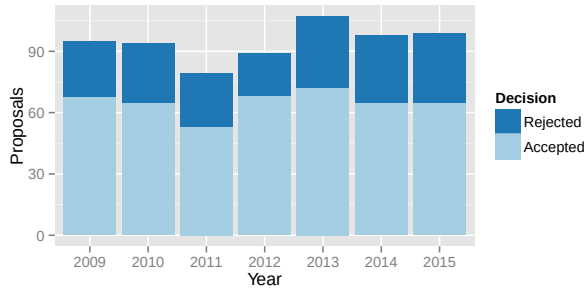

(a) Graphical distribution for 2009–2015

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
2009	95	68	71.6	27	28.4
2010	94	65	69.1	29	30.9
2011	79	53	67.1	26	32.9
2012	89	68	76.4	21	23.6
2013	107	72	67.3	35	32.7
2014	98	65	66.3	33	33.7
2015	99	65	65.7	34	34.3

(b) Detailed numbers for 2009–2015

Fig. 14.1

Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

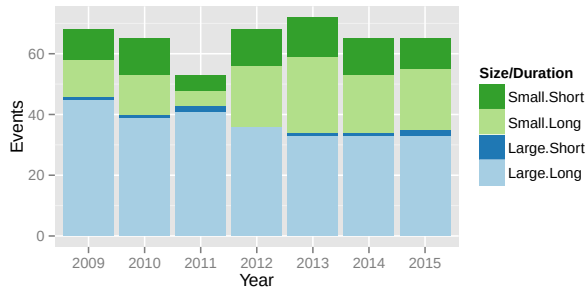

(a) Graphical distribution for 2009–2015

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2009	45	1	12	10	68
2010	39	1	13	12	65
2011	41	2	5	5	53
2012	36	0	20	12	68
2013	33	1	25	13	72
2014	33	1	19	12	65
2015	33	2	20	10	65

(b) Detailed numbers for 2009–2015

Fig. 14.2

Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2009–2015. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

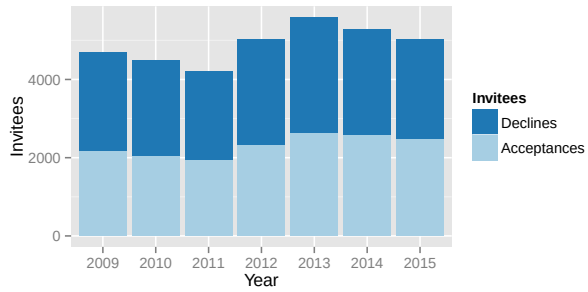

(a) Graphical distribution for 2009–2015

Year	Invitees	Accepances		Declines	
	#	#	%	#	%
2009	4694	2168	46.2	2526	53.8
2010	4499	2053	45.6	2446	54.4
2011	4223	1958	46.4	2265	53.6
2012	5033	2346	46.6	2687	53.4
2013	5591	2639	47.2	2952	52.8
2014	5285	2590	49.0	2695	51.0
2015	5023	2473	49.2	2550	50.8

(b) Detailed numbers for 2009–2015

Fig. 14.3

Invitees, accepting invitees, and declining invitees to Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

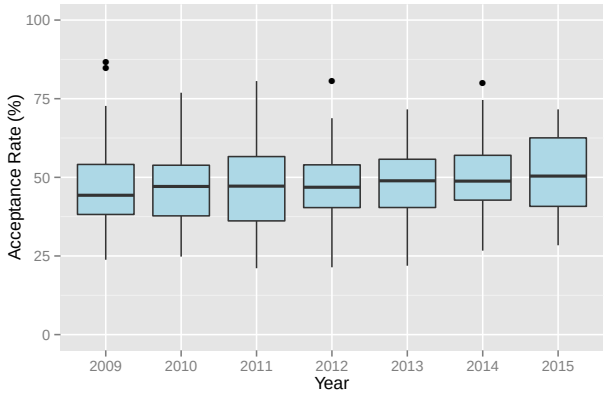

(a) Graphical distribution for 2009–2015

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2009	23.8	86.7	47.0	13.1
2010	24.8	76.9	46.5	12.2
2011	21.1	80.6	47.7	14.0
2012	21.4	80.5	47.2	11.0
2013	21.9	71.6	48.4	11.2
2014	26.7	80.0	50.2	11.2
2015	28.4	71.6	50.7	12.4

(b) Detailed numbers for 2009–2015

Fig. 14.4

Acceptance rate of guests invited to Dagstuhl Seminars or Dagstuhl Perspectives Workshops in 2009–2015. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

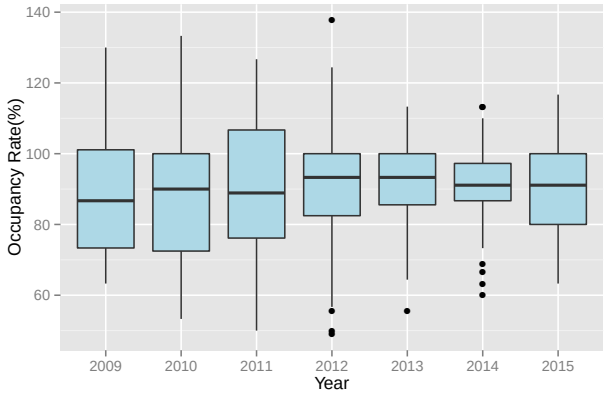

(a) Graphical distribution for 2009–2015

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2009	63.3	130.0	89.1	19.0
2010	53.3	133.3	85.9	18.0
2011	50.0	126.7	90.4	19.3
2012	48.9	137.8	92.4	17.6
2013	55.6	113.3	92.1	12.2
2014	60.0	113.3	90.6	10.3
2015	63.3	116.7	89.8	12.3

(b) Detailed numbers for 2009–2015

Fig. 14.5

Occupancy rate for Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2009–2015. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

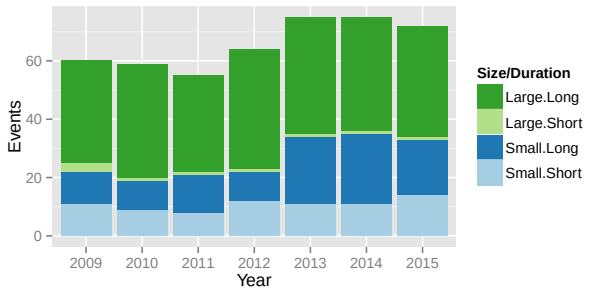

(a) Graphical distribution for 2009–2015

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2009	11	11	3	35	60
2010	9	10	1	39	59
2011	8	13	1	33	55
2012	12	10	1	41	64
2013	11	23	1	40	75
2014	11	24	1	39	75
2015	14	19	1	38	72

(b) Detailed numbers for 2009–2015

Fig. 14.6

Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2009–2015. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

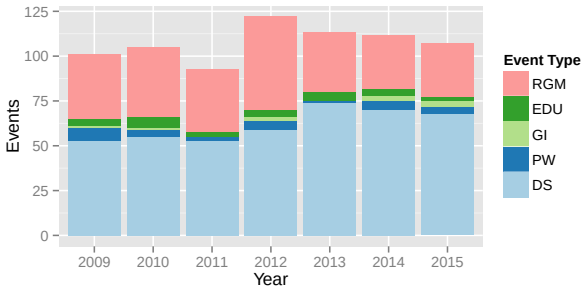

(a) Graphical distribution for 2009–2015

Year	DS	PW	GI	EDU	RGM	Total
2009	53	7	1	4	36	101
2010	55	4	1	6	39	105
2011	53	2	0	3	35	93
2012	59	5	2	4	52	122
2013	74	1	0	5	33	113
2014	70	5	3	4	30	112
2015	68	4	3	2	30	107

(b) Detailed numbers for 2009–2015

Fig. 14.7
Number of all events held at Dagstuhl, by type. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

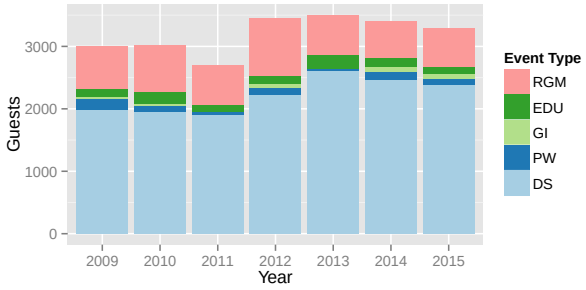

(a) Graphical distribution for 2009–2015

Year	DS		PW		GI		EDU		RGM		Total
	#	%	#	%	#	%	#	%	#	%	#
2009	1983	65.9	185	6.1	26	0.9	131	4.4	686	22.8	3011
2010	1950	64.7	103	3.4	25	0.8	192	6.4	743	24.7	3013
2011	1894	70.2	64	2.4	0	0.0	103	3.8	637	23.6	2698
2012	2226	64.4	120	3.5	48	1.4	144	4.2	916	26.5	3454
2013	2610	74.5	29	0.8	0	0.0	230	6.6	634	18.1	3503
2014	2463	72.2	127	3.7	86	2.5	144	4.2	589	17.3	3409
2015	2385	72.3	88	2.7	90	2.7	111	3.4	624	18.9	3298

(b) Detailed numbers for 2009–2015

Fig. 14.8
Number of participants. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

Country	A	B	Total
Germany	657	624	1281
United States	509	51	560
United Kingdom	257	7	264
France	158	9	167
Switzerland	78	10	88
Canada	77	7	84
Austria	72	11	83
Netherlands	59	14	73
Sweden	62	8	70
Italy	64	1	65
Luxembourg	17	38	55
Belgium	46	1	47
Spain	44	1	45
Japan	38	0	38
Finland	35	2	37
Australia	30	4	34
Denmark	30	2	32
Israel	30	2	32
India	21	2	23
China	16	2	18
Portugal	17	0	17
Russian Federation	16	1	17
Czech Republic	14	1	15
Poland	15	0	15
Pakistan	0	14	14
Brazil	13	0	13
Norway	10	1	11
Republic of Korea	8	2	10
Slovenia	9	1	10
Singapore	9	0	9
Hong Kong	8	0	8
Hungary	6	2	8
Greece	7	0	7
Ireland	6	0	6
Serbia	0	5	5
South Africa	5	0	5
Taiwan	5	0	5
New Zealand	4	0	4
Estonia	3	0	3
Malta	2	1	3
Turkey	3	0	3
Chile	2	0	2
Argentina	1	0	1
Croatia	1	0	1
Latvia	1	0	1
Mexico	1	0	1
Oman	1	0	1
Romania	1	0	1
Saudi Arabia	0	1	1
Slovak Republic	1	0	1
Ukraine	1	0	1
United Arab Emirates	1	0	1
Uruguay	1	0	1
Vietnam	1	0	1
Total	2473	825	3298

(a) Details for 2015 by country

Region	A		B		Total	
	#	%	#	%	#	%
Germany	657	26.6	624	75.6	1281	38.8
Europe (w/o Germany)	1035	41.9	116	14.1	1151	34.9
North America	586	23.7	58	7	644	19.5
Asia	138	5.6	23	2.8	161	4.9
Australia	34	1.4	4	0.5	38	1.2
South America	18	0.7	0	0	18	0.5
Africa	5	0.2	0	0	5	0.2
Total	2473	100	825	100	3298	100

(b) Details for 2015 by region

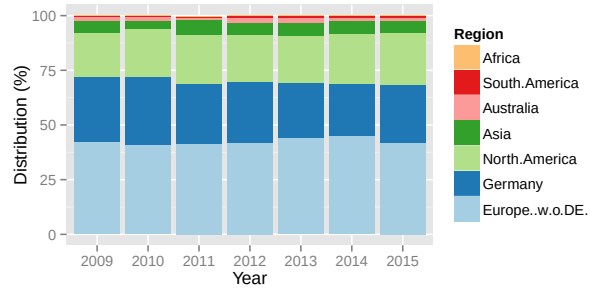

(c) Graphical distribution of seminar type A in 2009–2015 by year and region

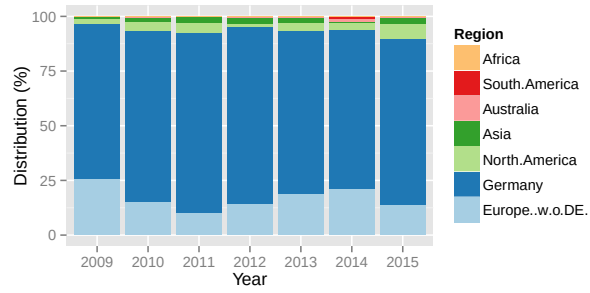

(d) Graphical distribution of seminar type B in 2009–2015 by year and region

Fig. 14.9

Number of Dagstuhl guests by country of origin. A = Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, B = Participants in all other events (GI-Dagstuhl Seminars, educational events, and research group meetings).

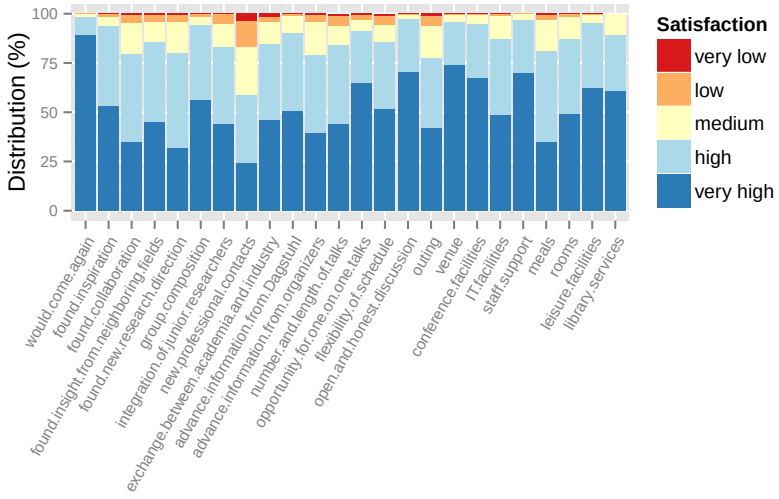

(a) Graphical distribution for 2015

	2009	2010	2011	2012	2013	2014	2015	2015 – 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	15	133	1295	1446
found inspiration	4.3	4.3	4.4	4.4	4.5	4.4	4.5	2	17	68	592	778	1457
found collaboration	3.9	3.9	4.0	4.0	4.0	4.1	4.1	10	52	232	639	509	1442
found insight from neighboring fields	4.1	4.0	4.1	4.1	4.1	4.2	4.3	9	46	151	591	657	1454
found new research direction	3.9	3.9	4.0	4.0	4.0	4.0	4.1	9	51	225	695	456	1436
group composition	4.4	4.4	4.4	4.4	4.4	4.4	4.5	3	20	57	558	824	1462
integration of junior researchers	4.2	4.2	4.2	4.2	4.2	4.2	4.2	3	69	173	564	641	1450
new professional contacts	3.5	3.5	3.7	3.6	3.7	3.7	3.6	53	181	349	484	347	1414
exchange between academia and industry	4.1	4.3	4.1	4.3	4.2	4.2	4.3	14	25	100	343	419	901
advance information from Dagstuhl	4.2	4.3	4.3	4.4	4.4	4.4	4.4	2	13	126	566	733	1440
advance information from organizers	4.1	4.0	4.2	4.1	4.1	4.1	4.1	10	46	241	555	557	1409
number and length of talks	4.0	4.1	4.1	4.1	4.2	4.1	4.2	14	74	140	583	637	1448
opportunity for one on one talks	4.4	4.5	4.5	4.4	4.5	4.5	4.5	8	39	76	387	933	1443
flexibility of schedule	4.2	4.3	4.2	4.2	4.2	4.3	4.3	12	71	123	489	744	1439
open and honest discussion	4.7	4.7	4.7	4.6	4.7	4.7	4.7	4	6	25	393	1022	1450
outing	4.1	4.0	4.2	4.1	4.1	4.1	4.1	13	51	181	386	457	1088
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.7	3	7	46	318	1082	1456
conference facilities	4.8	4.7	4.8	4.7	4.6	4.7	4.6	2	4	68	398	983	1455
IT facilities	4.5	4.5	4.6	4.4	4.4	4.4	4.3	3	12	150	499	633	1297
staff support	4.7	4.7	4.7	4.7	4.7	4.7	4.7	0	2	40	374	976	1392
meals	4.3	4.2	4.3	4.2	4.1	4.1	4.1	9	32	233	670	510	1454
rooms	4.5	4.5	4.5	4.4	4.4	4.4	4.4	3	17	164	548	714	1446
leisure facilities	4.5	4.5	4.6	4.6	4.6	4.6	4.6	1	3	59	442	847	1352
library services	4.6	4.6	4.5	4.5	4.5	4.5	4.5	0	1	68	189	397	655

(b) Averages for 2009–2015 and detailed numbers for 2015: 1 = very low, 2 = low, 3 = medium, 4 = high, 5 = very high

Fig. 14.10
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

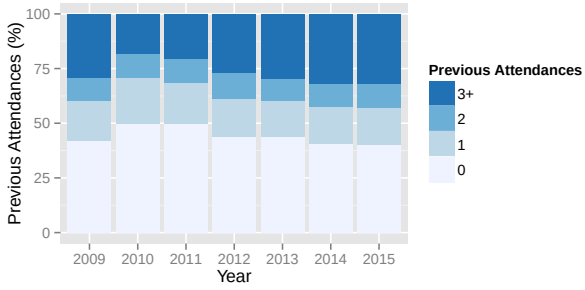

(a) Graphical distribution for 2009–2015

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	#
2009	446	42	193	18	114	11	307	29	1060
2010	442	50	185	21	98	11	162	18	887
2011	413	50	154	19	94	11	168	20	829
2012	483	44	193	17	135	12	295	27	1106
2013	630	44	237	17	145	10	422	29	1434
2014	561	40	239	17	144	10	443	32	1387
2015	573	40	234	17	158	11	451	32	1416

(b) Detailed numbers for 2009–2015

Fig. 14.11

Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants and their previous instances of attendance in Dagstuhl Seminars or Dagstuhl Perspectives Workshops, according to our guest survey.

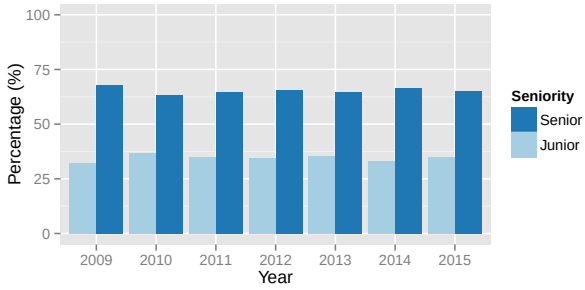

(a) Graphical distribution for 2009–2015

Year	Junior		Senior		Total
	#	%	#	%	#
2009	269	32.3	565	67.7	834
2010	291	36.8	500	63.2	791
2011	266	35.2	489	64.8	755
2012	307	34.6	580	65.4	887
2013	413	35.4	754	64.6	1167
2014	382	33.3	765	66.7	1147
2015	410	34.9	764	65.1	1174

(b) Detailed numbers for 2009–2015

Fig. 14.12

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants, according to our guest survey.

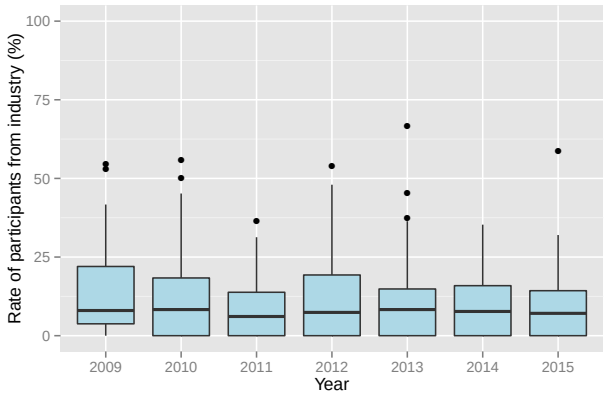

(a) Graphical distribution for 2009–2015

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2009	0.0	54.5	13.6	13.8
2010	0.0	56.0	12.2	13.1
2011	0.0	36.4	8.2	9.1
2012	0.0	53.8	12.4	13.7
2013	0.0	66.7	11.6	12.8
2014	0.0	35.3	9.4	9.4
2015	0.0	58.8	9.8	10.5

(b) Detailed numbers for 2009–2015

Fig. 14.13

Distribution of self-assigned primary occupation in business of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2009–2015, 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”.

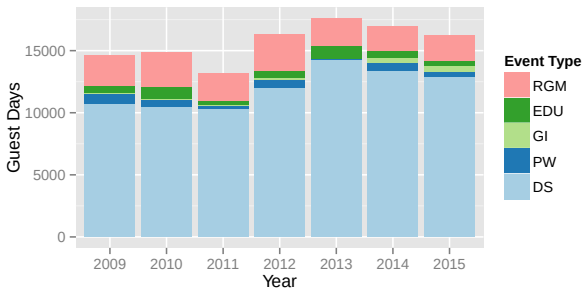

(a) Graphical distribution for 2009–2015

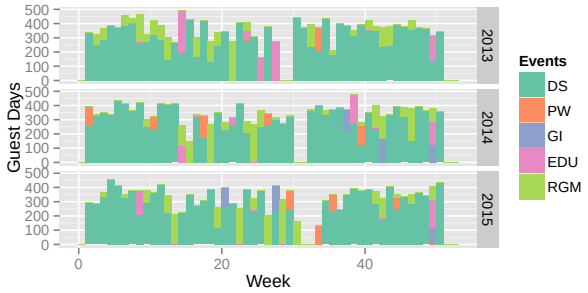

(b) Graphical distribution for 2013–2015 by week

Year	DS	PW	GI	EDU	RGM	Total
2009	10700	842	103	509	2462	14616
2010	10522	484	150	914	2745	14815
2011	10309	292	0	369	2241	13211
2012	12024	578	238	537	2947	16324
2013	14222	159	0	983	2248	17612
2014	13402	602	434	534	1959	16931
2015	12876	468	434	372	2048	16198

(c) Detailed numbers for 2009–2015

Fig. 14.14

Number of guest days at Schloss Dagstuhl. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, RGM = research group meeting.

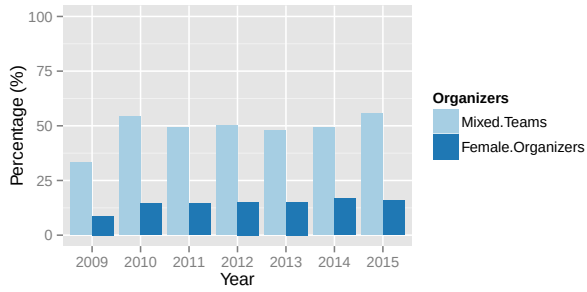

(a) Graphical distribution for 2009–2015

	Organizer Teams			Organizers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2009	60	20	33.3	229	20	8.7
2010	59	32	54.2	233	34	14.6
2011	55	27	49.1	213	31	14.6
2012	64	32	50.0	256	39	15.2
2013	75	36	48.0	282	43	15.2
2014	75	37	49.3	303	51	16.8
2015	72	40	55.6	284	45	15.8

(b) Detailed numbers for 2009–2015

Fig. 14.15

Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

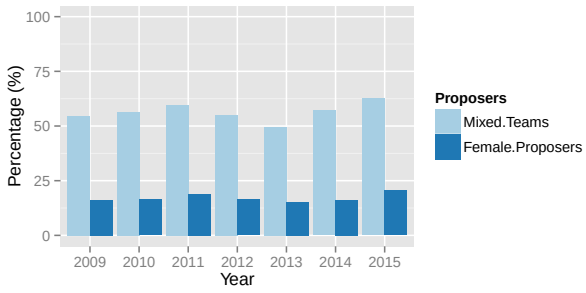

(a) Graphical distribution for 2009–2015

	Proposer Teams			Proposers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2009	95	52	54.7	365	59	16.2
2010	94	53	56.4	366	60	16.4
2011	79	47	59.5	311	59	19.0
2012	89	49	55.1	341	56	16.4
2013	107	53	49.5	431	66	15.3
2014	98	56	57.1	387	63	16.3
2015	99	62	62.6	391	80	20.5

(b) Detailed numbers for 2009–2015

Fig. 14.16
Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals with mixed-gender proposer teams.

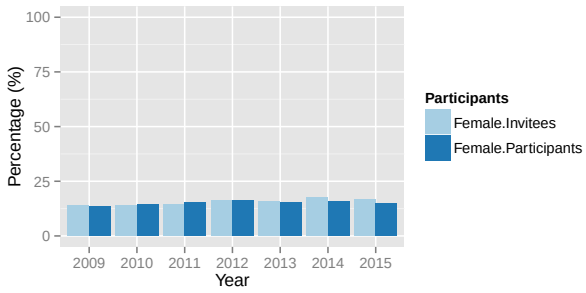

(a) Graphical distribution for 2009–2015

	Invitees			Participants		
Year	Total	Female		Total	Female	
	#	#	%	#	#	%
2009	4694	652	13.9	2168	296	13.7
2010	4499	632	14.0	2053	294	14.3
2011	4223	604	14.3	1958	295	15.1
2012	5033	822	16.3	2346	377	16.1
2013	5591	889	15.9	2639	401	15.2
2014	5285	943	17.8	2590	406	15.7
2015	5023	845	16.8	2473	369	14.9

(b) Detailed numbers for 2009–2015

Fig. 14.17
Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.

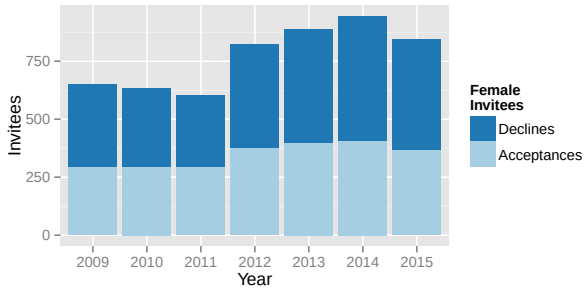

(a) Graphical distribution for 2009–2015

Year	Female Invitees	Acceptances		Declines	
	#	#	%	#	%
2009	652	296	45.4	356	54.6
2010	632	294	46.5	338	53.5
2011	604	295	48.8	309	51.2
2012	822	377	45.9	445	54.1
2013	889	401	45.1	488	54.9
2014	943	406	43.1	537	56.9
2015	845	369	43.7	476	56.3

(b) Detailed numbers for 2009–2015

Fig. 14.18
Female invitees to Dagstuhl Seminar and Dagstuhl Perspectives Workshops.

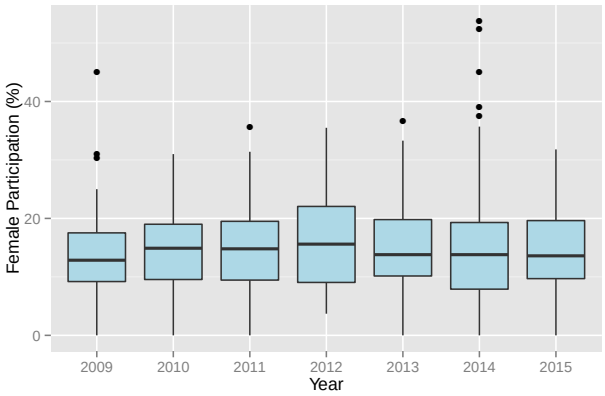

(a) Graphical distribution for 2009–2015

Year	Min (%)	Max (%)	Avg (%)	Std (%)
2009	0.0	45.0	13.8	7.7
2010	0.0	31.0	14.5	7.7
2011	0.0	35.7	14.7	7.4
2012	3.7	35.5	16.1	7.7
2013	0.0	36.7	15.1	7.3
2014	0.0	53.8	15.9	11.1
2015	0.0	31.8	14.8	7.7

(b) Detailed numbers for 2009–2015

Fig. 14.19

Distribution of female participants rate of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2009–2015. Min = minimal value, Max = maximal value, Avg = average, Std = standard deviation.

Statistiken zur Bibliographiedatenbank dblp

14.2

Statistics of the dblp computer science bibliography

Dieser Abschnitt enthält statistische Daten zur Bibliographiedatenbank dblp. Fig. 14.20 listet die durchschnittlichen Nutzungszahlen der letzten Jahre. Ein Überblick über die Entwicklung des dblp Datenbestandes kann Fig. 14.21 und Fig. 14.22 entnommen werden.

This section provides statistical data about the dblp computer science bibliography. Fig. 14.20 show the average usage statistics of the dblp servers of the past years. An overview of the development of the dblp database can be found in Fig. 14.21 and Fig. 14.22.

	Trier 1		Trier 2		Dagstuhl	
	2014	2015	2014	2015	2014	2015
user sessions (visits) per day	21,057	28,327	4,703	662	326	510
page views per day	174,247	452,089	47,531	8,839	14,964	14,868
page views per user session	8.2	15.9	10.1	13.3	45.8	29.1
distinct users (IPs) per month	327,299	416,413	76,566	11,474	4,399	7,241
data served per month	825.2 GB	861.8 GB	345.8 GB	22.0 GB	27.2 GB	75.8 GB
as above, including bots	5,187.5 GB	2,206.6 GB	672.2 GB	69.5 GB	56.4 GB	172.3 GB

Fig. 14.20

Average usage of the three dblp servers. Trier 1 = http://dblp.uni-trier.de, Trier 2 = http://dblp2.uni-trier.de, Dagstuhl = http://dblp.dagstuhl.de. In 2015, changes have been made in the server setup in order to shift traffic from development server Trier 2 to the more powerful server Trier 1.

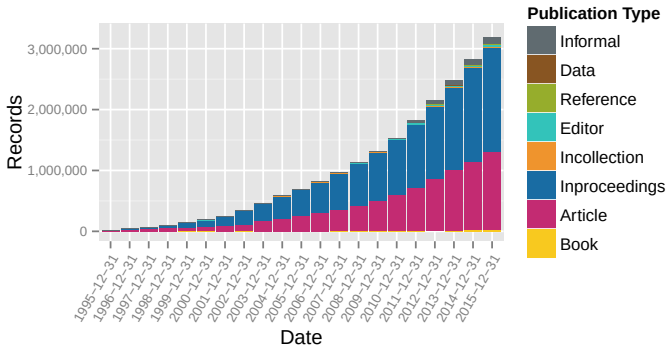

(a) Graphical distribution for 1996–2015

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2009	1,345	0.1	506,778	38.6	778,060	59.2	14,522	1.1	12,967	1.0	0	0.0	47	0.0	1,313,719
2010	1,430	0.1	604,329	39.7	887,068	58.3	15,113	1.0	14,490	1.0	0	0.0	47	0.0	1,522,477
2011	9,343	0.5	706,467	38.7	1,038,724	56.9	7,451	0.4	16,869	0.9	12,201	0.7	35,796	2.0	1,826,851
2012	16,023	0.7	847,659	39.4	1,189,533	55.3	9,623	0.4	19,955	0.9	13,123	0.6	53,251	2.5	2,149,167
2013	16,814	0.7	995,550	40.2	1,342,872	54.3	12,790	0.5	22,606	0.9	13,124	0.5	69,847	2.8	2,473,603
2014	17,531	0.6	1,126,808	39.9	1,539,151	54.4	14,453	0.5	26,036	0.9	14,689	0.5	88,125	3.1	2,826,793
2015	18,313	0.6	1,278,941	40.0	1,721,215	53.9	16,260	0.5	30,012	0.9	19,103	0.6	110,813	3.5	3,194,657

(b) Detailed numbers for 2009–2015

Fig. 14.21
Development of the total size of the dblp database.

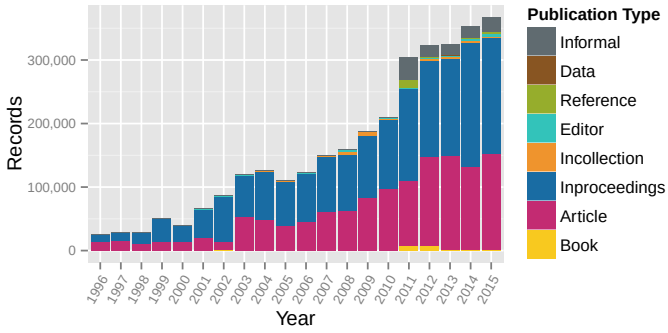

(a) Graphical distribution for 1996–2015

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2009	92	0.0	83,635	44.5	96,679	51.4	6,184	3.3	1,537	0.8	0	0.0	0	0.0	188,127
2010	85	0.0	97,551	46.7	109,008	52.2	591	0.3	1,523	0.7	0	0.0	0	0.0	208,758
2011	7,913	2.6	102,138	33.6	151,656	49.8	−7,662	−2.5	2,379	0.8	12,201	4.0	35,749	11.7	304,374
2012	6,680	2.1	141,192	43.8	150,809	46.8	2,172	0.7	3,086	1.0	922	0.3	17,455	5.4	322,316
2013	791	0.2	147,891	45.6	153,339	47.3	3,167	1.0	2,651	0.8	1	0.0	16,596	5.1	324,436
2014	717	0.2	131,258	37.2	196,279	55.6	1,663	0.5	3,430	1.0	1,565	0.4	18,278	5.2	353,190
2015	782	0.2	152,133	41.4	182,064	49.5	1,807	0.5	3,976	1.1	4,414	1.2	22,688	6.2	367,864

(b) Detailed numbers for 2009–2015

Fig. 14.22
Rate of inclusion of new publications to dblp. The negative number of new *Incollection* records in 2011 results from relabeling several thousand existing records with the newly introduced *Reference* type. Similarily, in the same year, several thousand *Article* and *Inproceedings* records have been labeled as *Informal*.

Statistiken zu Dagstuhl Publishing

14.3 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. 14.23 fasst die statistischen Daten der Veröffentlichungen in der Zeitschrift Dagstuhl Reports zusammen, Fig. 14.24 die der Publikationen in der Reihe Dagstuhl Manifestos und schließlich Fig. 14.25 die der veröffentlichten Bände in der Reihe Dagstuhl Follow-Ups.

Die statistischen Daten zu den dienstleistungsbezogenen Veröffentlichungen finden sich anschließend: Fig. 14.26 fasst die Daten in der Reihe OASIcs und Fig. 14.27 die der Reihe LIPIcs zusammen.

Fig. 14.28 fasst die Kennzahlen der Zeitschrift LITES zusammen.

Die Abbildungen in Fig. 14.30 und 14.31 zeigen die Anzahl von Zugriffen, die seit 2016 standardisiert über den *Open Access Statistik-Service*[57] der GBV Göttingen ausgewertet werden.

Die verschiedenen Publikationsserien wurden in verschiedenen Jahren zwischen 2009 und 2015 gegründet. Wir stellen in den Statistiken dennoch stets den gesamten Zeitraum dar.

In this section the statistical data of Dagstuhl Publishing are presented.

The first three figures present the development of the seminar-focussed series: Fig. 14.23 summarizes the data of the periodical Dagstuhl Reports, Fig. 14.24 the data of the Dagstuhl Manifestos series, and, finally, Fig. 14.25 those of the volumes published in the Dagstuhl Follow-Ups series.

The statistical data to the service-focussed series are presented afterwards. Fig. 14.26 presents numbers related to OASIcs and Fig. 14.27 numbers related to LIPIcs.

We summarize the publications of the journal LITES in Fig. 14.28.

The figures in Fig. 14.30 and 14.31 show the number of online accesses, which are monitored since 2016 in a standardized procedure using the *Open Access Statistik-Service*[57] offered by GBV Göttingen.

Please note that the publication series were established in different years, but all in the period between 2009 and 2015. However, we always consider this complete period.

[57] https://www.gbv.de/Verbundzentrale/serviceangebote/oas-service/open-access-statistik-service

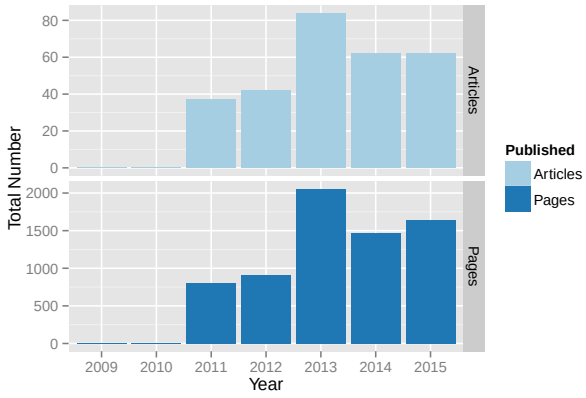

(a) Graphical distribution for 2011–2015

Year	Articles	Pages
2009	0	0
2010	0	0
2011	37	806
2012	42	913
2013	84	2059
2014	62	1464
2015	62	1636

(b) Detailed numbers for 2011–2015

Fig. 14.23
Statistics about Dagstuhl Reports published between 2011 to 2015.

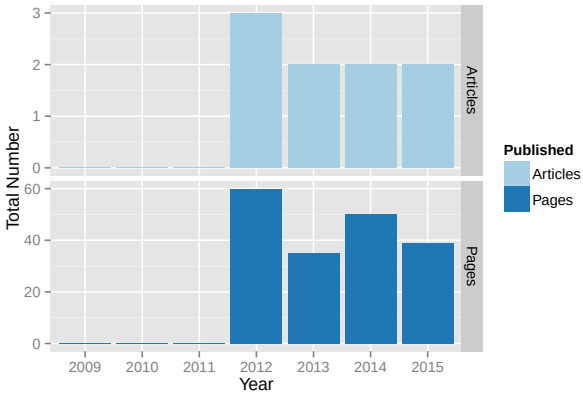

(a) Graphical distribution for 2012–2015

Year	Articles	Pages
2009	0	0
2010	0	0
2011	0	0
2012	3	60
2013	2	35
2014	2	50
2015	2	39

(b) Detailed numbers for 2012–2015

Fig. 14.24
Statistics about Dagstuhl Manifestos published between 2012 to 2015.

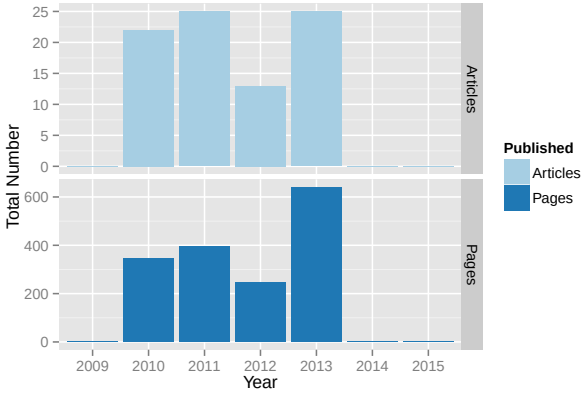

(a) Graphical distribution for 2010–2015

Year	Volumes	Articles	Pages
2009	0	0	0
2010	1	22	345
2011	1	25	395
2012	1	13	246
2013	3	25	641
2014	0	0	0
2015	0	0	0

(b) Detailed numbers for 2010–2015

Fig. 14.25
Statistics about Dagstuhl Follow-Ups volumes published between 2010 to 2015.

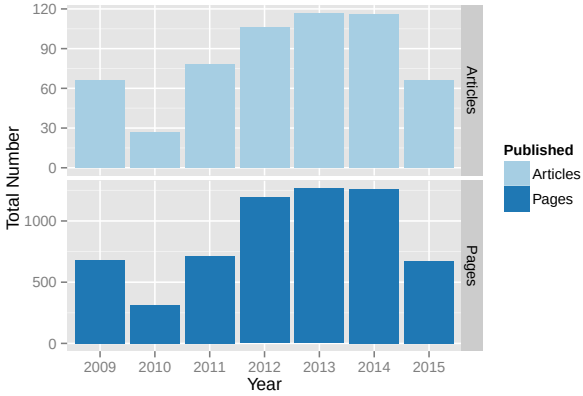

(a) Graphical distribution for 2009–2015

Year	Volumes	Articles	Pages
2009	4	66	679
2010	2	27	315
2011	5	78	717
2012	8	106	1192
2013	7	117	1265
2014	8	116	1264
2015	6	66	674

(b) Detailed numbers for 2009–2015

Fig. 14.26
Statistics about OASIcs volumes published between 2009 to 2015.

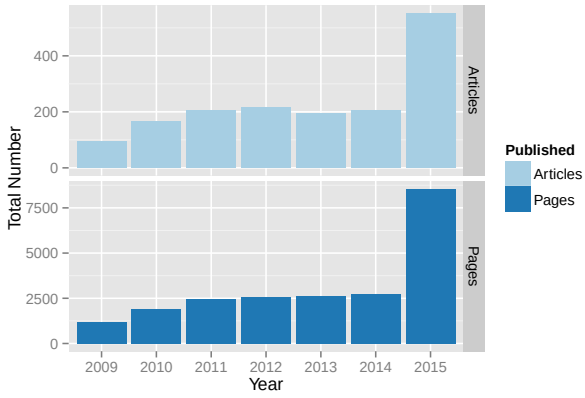

(a) Graphical distribution for 2009–2015

Year	Volumes	Articles	Pages
2009	2	95	1169
2010	4	167	1907
2011	5	205	2439
2012	5	215	2591
2013	6	195	2607
2014	5	204	2752
2015	16	553	8565

(b) Detailed numbers for 2009–2015

Fig. 14.27
Statistics about LIPIcs volumes published between 2009 to 2015.

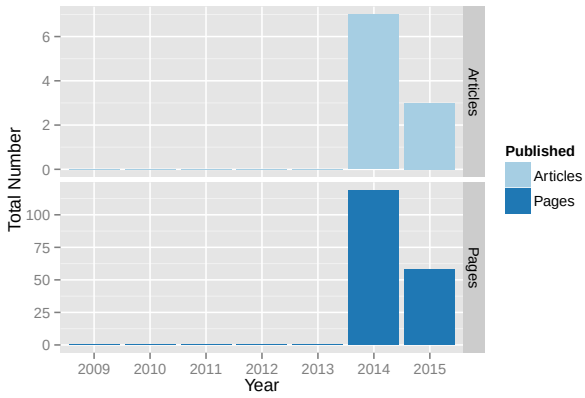

(a) Graphical distribution for 2014–2015

Year	Articles	Pages
2009	0	0
2010	0	0
2011	0	0
2012	0	0
2013	0	0
2014	7	119
2015	3	58

(b) Detailed numbers for 2014–2015

Fig. 14.28
Statistics about LITES articles published between 2014 to 2015.

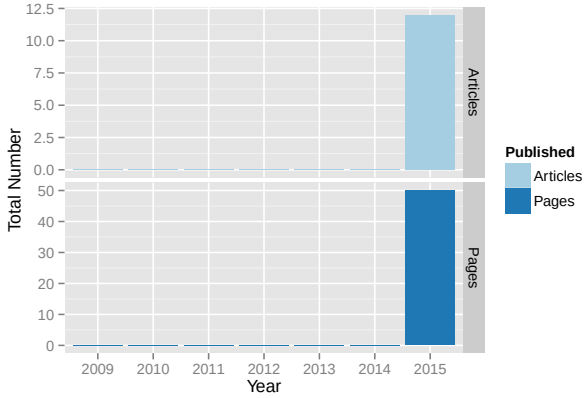

(a) Graphical distribution for 2011–2015

Year	Articles	Pages
2009	0	0
2010	0	0
2011	0	0
2012	0	0
2013	0	0
2014	0	0
2015	12	50

(b) Detailed numbers for 2011–2015

Fig. 14.29
Statistics about DARTS artifacts published between 2011 to 2015.

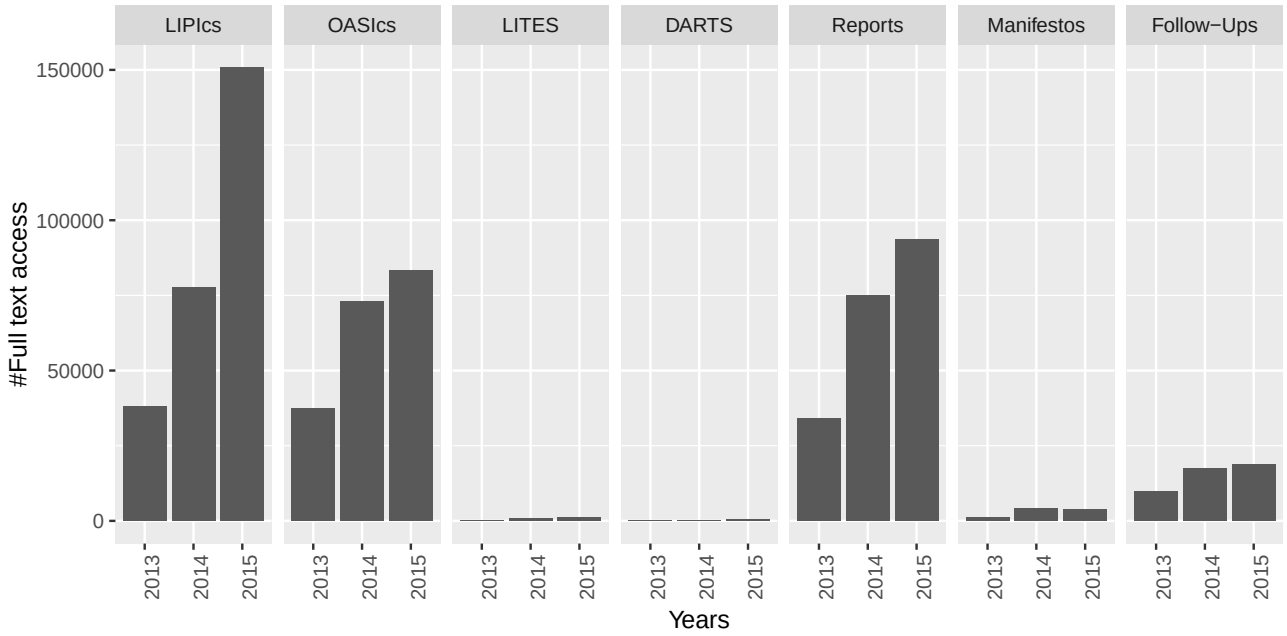

Fig. 14.30
Total number of full text accesses for articles published between 2013 to 2015.

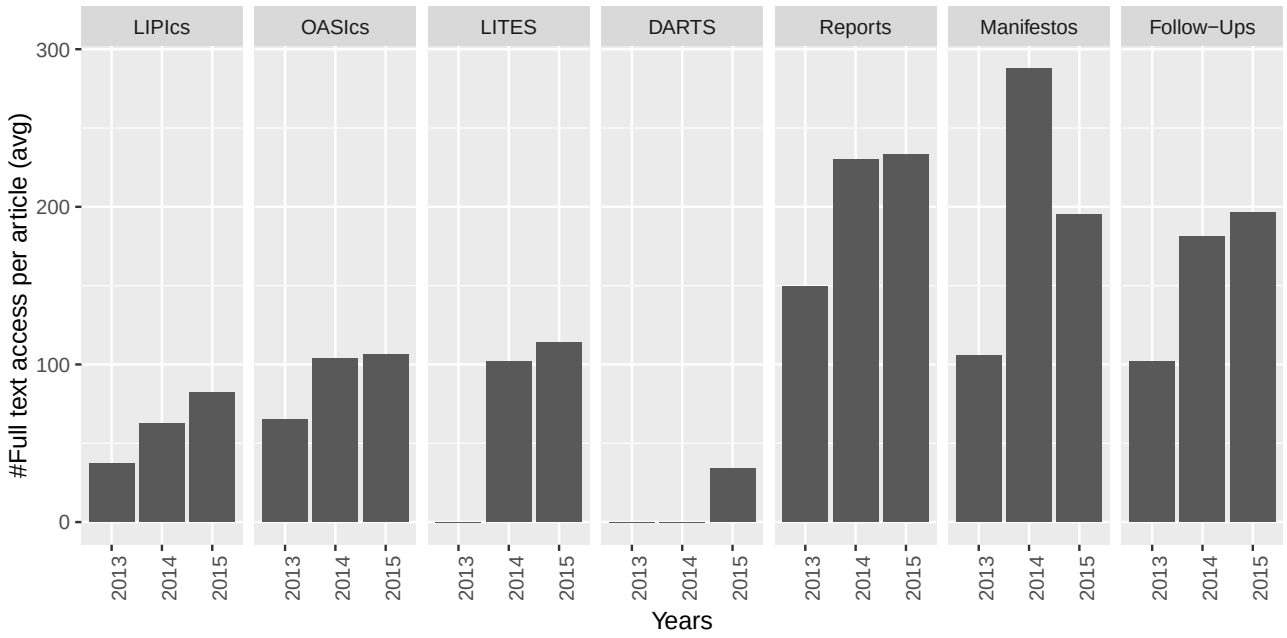

Fig. 14.31
Average number of full text accesses per article for articles published between 2013 to 2015.

15 Veranstaltungen 2015
Schedule of Events 2015

Dagstuhl-Seminare 15.1 *Dagstuhl Seminars*

15021 – Concurrent Computing in the Many-Core Era
Pascal Felber (Université de Neuchâtel, CH), J. Eliot B. Moss (University of Massachusetts – Amherst, US), Michael Philippsen (Universität Erlangen-Nürnberg, DE), Michael L. Scott (University of Rochester, US)
January 4–9, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15021

15022 – Quality of Experience: From Assessment to Application
Katrien De Moor (NTNU – Trondheim, NO), Markus Fiedler (Blekinge Institute of Technology – Karlskrona, SE), Peter Reichl (Universität Wien, AT), Martín Varela (VTT Technical Research Centre of Finland – Oulu, FI)
January 4–7, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15022

15031 – Understanding Complexity in Multiobjective Optimization
Salvatore Greco (University of Catania, IT & University of Portsmouth, GB), Kathrin Klamroth (Universität Wuppertal, DE), Joshua D. Knowles (University of Manchester, GB), Günter Rudolph (TU Dortmund, DE)
January 11–16, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15031

15041 – Model-Driven Algorithms and Architectures for Self-Aware Computing Systems
Jeffrey O. Kephart (IBM TJ Watson Research Center – Yorktown Heights, US), Samuel Kounev (Universität Würzburg, DE), Marta Kwiatkowska (University of Oxford, GB), Xiaoyun Zhu (VMware – Palo Alto, US)
January 18–23, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15041

15042 – Coalgebraic Semantics of Reflexive Economics
Samson Abramsky (University of Oxford, GB), Alexander Kurz (University of Leicester, GB), Pierre Lescanne (ENS – Lyon, FR), Viktor Winschel (Universität Mannheim, DE)
January 18–21, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15042

15051 – Artificial and Computational Intelligence in Games: Integration
Simon M. Lucas (University of Essex, GB), Michael Mateas (University of California – Santa Cruz, US), Mike Preuß (Universität Münster, DE), Pieter Spronck (Tilburg University, NL), Julian Togelius (New York University, US)
January 25–30, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15051

15052 – Empirical Evaluation for Graph Drawing
Ulrik Brandes (Universität Konstanz, DE), Irene Finocchi (Sapienza University of Rome, IT), Martin Nöllenburg (KIT – Karlsruher Institut für Technologie, DE), Aaron Quigley (University of St. Andrews, GB)
January 25–30, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15052

15061 – Non-Zero-Sum-Games and Control
Krishnendu Chatterjee (IST Austria – Klosterneuburg, AT), Stéphane Lafortune (University of Michigan – Ann Arbor, US), Nicolas Markey (ENS – Cachan, FR), Wolfgang Thomas (RWTH Aachen, DE)
February 1–6, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15061

15062 – Domain-Specific Languages
Sebastian Erdweg (TU Darmstadt, DE), Martin Erwig (Oregon State University, US), Richard F. Paige (University of York, GB), Eelco Visser (TU Delft, NL)
February 1–6, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15062

15071 – Formal Foundations for Networking
Nikolaj S. Bjorner (Microsoft Corporation – Redmond, US), Nate Foster (Cornell University, US), Philip Brighten Godfrey (University of Illinois – Urbana-Champaign, US), Pamela Zave (AT&T Labs Research – Bedminster, US)
February 8–13, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15071

15072 – Distributed Cloud Computing
Yvonne Coady (University of Victoria, CA), James Kempf (Ericsson – San Jose, US), Rick McGeer (HP Enterprise Services – Palo Alto, US), Stefan Schmid (TU Berlin, DE)
February 8–11, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15072

15081 – Holistic Scene Understanding
Laura Leal-Taixé (ETH Zürich, CH), Jiri Matas (Czech Technical University, CZ), Vittorio Murino (Italian Institute of Technology – Genova, IT), Bodo Rosenhahn (Leibniz Universität Hannover, DE)
February 15–20, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15081

15082 – Limitations of Convex Programming: Lower Bounds on Extended Formulations and Factorization Ranks
Hartmut Klauck (Nanyang TU – Singapore, SG), Troy Lee (National University of Singapore, SG), Dirk Oliver Theis (University of Tartu, EE), Rekha R. Thomas (University of Washington – Seattle, US)
February 15–20, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15082

15091 – Smart Buildings and Smart Grids
Christoph Goebel (TU München, DE), Hans-Arno Jacobsen (TU München, DE), Randy H. Katz (University of California – Berkeley, US), Hartmut Schmeck (KIT – Karlsruher Institut für Technologie, DE)
February 22–27, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15091

15101 – Bridging Information Visualization with Machine Learning
Daniel A. Keim (Universität Konstanz, DE), Tamara Munzner (University of British Columbia – Vancouver, CA), Fabrice Rossi (University of Paris I, FR), Michel Verleysen (University of Louvain, BE)
March 1–6, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15101

15102 – Secure Routing for Future Communication Networks
Amir Herzberg (Bar-Ilan University – Ramat Gan, IL), Matthias Hollick (TU Darmstadt, DE), Allison Mankin (Verisign Labs – Reston, US), Adrian Perrig (ETH Zürich, CH)
March 1–4, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15102

15111 – Computational Geometry
Otfried Cheong (KAIST – Daejeon, KR), Jeff Erickson (University of Illinois – Urbana-Champaign, US), Monique Teillaud (INRIA Lorraine – Nancy, FR)
March 8–13, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15111

15112 – Network Calculus
Florin Ciucu (University of Warwick – Coventry, GB), Markus Fidler (Leibniz Universität Hannover, DE), Jörg Liebeherr (University of Toronto, CA), Jens Schmitt (TU Kaiserslautern, DE)
March 8–11, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15112

15121 – Mixed Criticality on Multicore/Manycore Platforms
Sanjoy K. Baruah (University of North Carolina at Chapel Hill, US), Liliana Cucu-Grosjean (INRIA – Le Chesnay, FR), Robert Davis (University of York, GB), Claire Maiza (VERIMAG – Grenoble, FR)
March 15–20, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15121

15122 – Formal Models of Graph Transformation in Natural Language Processing
Frank Drewes (University of Umeå, SE), Kevin Knight (USC – Marina del Rey, US), Marco Kuhlmann (Linköping University, SE)
March 15–20, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15122

15131 – Normative Multi-Agent Systems
Amit K. Chopra (Lancaster University, GB), Leon van der Torre (University of Luxembourg, LU), Harko Verhagen (Stockholm University, SE), Serena Villata (INRIA Sophia Antipolis – Méditerranée, FR)
March 22–27, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15131

15151 – Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations
David Hutchison (Lancaster University, GB), Klara Nahrstedt (University of Illinois – Urbana-Champaign, US), Marcus Schöller (Hochschule Reutlingen, DE), Indra Spiecker gen. Döhmann (Goethe-Universität Frankfurt, DE), Markus Tauber (AIT Austrian Institute of Technology – Wien, AT)
April 7–10, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15151

15152 – Machine Learning with Interdependent and Non-Identically Distributed Data
Trevor Darrell (University of California – Berkeley, US), Marius Kloft (HU Berlin, DE), Massimiliano Pontil (University College London, GB), Gunnar Rätsch (Memorial Sloan-Kettering Cancer Center – New York, US), Erik Rodner (Universität Jena, DE)
April 7–10, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15152

15161 – Advanced Stencil-Code Engineering
Matthias Bolten (Universität Wuppertal, DE), Robert D. Falgout (LLNL – Livermore, US), Christian Lengauer (Universität Passau, DE), Olaf Schenk (University of Lugano, CH)
April 12–17, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15161

15162 – Software and Systems Traceability for Safety-Critical Projects
Jane Cleland-Huang (DePaul University – Chicago, US), Patrick Mäder (TU Ilmenau, DE), Sanjai Rayadurgam (University of Minnesota – Minneapolis, US), Wilhelm Schäfer (Universität Paderborn, DE)
April 12–17, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15162

15171 – Theory and Practice of SAT Solving
Armin Biere (Universität Linz, AT), Vijay Ganesh (University of Waterloo, CA), Martin Grohe (RWTH Aachen, DE), Jakob Nordström (KTH Royal Institute of Technology, SE), Ryan Williams (Stanford University, US)
April 19–24, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15171

15181 – Challenges and Trends in Probabilistic Programming
Gilles Barthe (IMDEA Software – Madrid, ES), Andrew D. Gordon (Microsoft Research UK – Cambridge, GB), Joost-Pieter Katoen (RWTH Aachen, DE), Annabelle McIver (Macquarie University – Sydney, AU)
April 26–30, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15181

15182 – Qualification of Formal Methods Tools
Darren Cofer (Rockwell Collins – Bloomington, US), Gerwin Klein (Data61 / NICTA – Sydney, AU), Konrad Slind (Rockwell Collins – Bloomington, US), Virginie Wiels (ONERA – Toulouse, FR)
April 26–29, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15182

15191 – Compositional Verification Methods for Next-Generation Concurrency
Lars Birkedal (Aarhus University, DK), Derek Dreyer (MPI-SWS – Saarbrücken, DE), Philippa Gardner (Imperial College London, GB), Zhong Shao (Yale University, US)
May 3–8, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15191

15192 – The Message in the Shadow: Noise or Knowledge?
Roberto Casati (ENS – Paris, FR), Patrick Cavanagh (Paris Descartes University, FR), Paulo E. Santos (University Center of FEI – Sao Paolo, BR)
May 3–8, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15192

15201 – Cross-Lingual Cross-Media Content Linking: Annotations and Joint Representations
Alexander G. Hauptmann (Carnegie Mellon University, US), James Hodson (Bloomberg – New York, US), Juanzi Li (Tsinghua University – Beijing, CN), Achim Rettinger (KIT – Karlsruher Institut für Technologie, DE), Nicu Sebe (University of Trento, IT)
May 10–13, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15201

15211 – Theory of Evolutionary Algorithms
Benjamin Doerr (Ecole Polytechnique – Palaiseau, FR), Nikolaus Hansen (University of Paris South XI, FR), Christian Igel (University of Copenhagen, DK), Lothar Thiele (ETH Zürich, CH)
May 17–22, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15211

15221 – Multi-Disciplinary Approaches to Reasoning with Imperfect Information and Knowledge – a Synthesis and a Roadmap of Challenges
Igor Douven (Paris-Sorbonne University, FR), Gabriele Kern-Isberner (TU Dortmund, DE), Markus Knauff (Universität Gießen, DE), Henri Prade (Paul Sabatier University – Toulouse, FR)
May 25–29, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15221

15222 – Human-Centric Development of Software Tools
Andrew J. Ko (University of Washington – Seattle, US), Shriram Krishnamurthi (Brown University – Providence, US), Gail C. Murphy (University of British Columbia – Vancouver, CA), Janet Siegmund (Universität Passau, DE)
May 25–28, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15222

15241 – Computational Social Choice: Theory and Applications
Craig Boutilier (University of Toronto, CA), Britta Dorn (Universität Tübingen, DE), Nicolas Maudet (UPMC – Paris, FR), Vincent Merlin (Caen University, FR)
June 7–12, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15241

15242 – Complexity of Symbolic and Numerical Problems
Peter Bürgisser (TU Berlin, DE), Felipe Cucker (City University – Hong Kong, HK), Marek Karpinski (Universität Bonn, DE), Nicolai Vorobjov (University of Bath, GB)
June 7–12, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15242

15251 – Sparse Modelling and Multi-Exponential Analysis
Annie Cuyt (University of Antwerp, BE), George Labahn (University of Waterloo, CA), Wen-Shin Lee (University of Antwerp, BE), Avraham Sidi (Technion – Haifa, IL)
June 14–19, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15251

15261 – Logics for Dependence and Independence
Erich Grädel (RWTH Aachen, DE), Juha Kontinen (University of Helsinki, FI), Jouko Väänänen (University of Helsinki, FI), Heribert Vollmer (Leibniz Universität Hannover, DE)
June 21–26, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15261

15262 – Life-Long Health Behavior-Change Technologies
Susanne Boll (Universität Oldenburg, DE), Eric Hekler (Arizona State University – Phoenix, US), Predrag Klasnja (University of Michigan – Ann Arbor, US)
June 21–26, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15262

15281 – Algorithms and Scheduling Techniques to Manage Resilience and Power Consumption in Distributed Systems
Henri Casanova (University of Hawaii at Manoa – Honolulu, US), Ewa Deelman (USC – Marina del Rey, US), Yves Robert (ENS – Lyon, FR), Uwe Schwiegelshohn (TU Dortmund, DE)
July 5–10, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15281

15301 – The Constraint Satisfaction Problem: Complexity and Approximability
Andrei A. Bulatov (Simon Fraser University – Burnaby, CA), Venkatesan Guruswami (Carnegie Mellon University, US), Andrei Krokhin (Durham University, GB), Dániel Marx (Hungarian Academy of Sciences – Budapest, HU)
July 19–24, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15301

15351 – Computational Mass Spectrometry
Rudolf Aebersold (ETH Zürich, CH), Oliver Kohlbacher (Universität Tübingen, DE), Olga Vitek (Northeastern University – Boston, US)
August 23–28, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15351

15352 – Design of Microfluidic Biochips: Connecting Algorithms and Foundations of Chip Design to Biochemistry and the Life Sciences
Krishnendu Chakrabarty (Duke University – Durham, US), Tsung-Yi Ho (National Chiao-Tung University – Hsinchu, TW), Robert Wille (Universität Bremen, DE)
August 23–26, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15352

15361 – Mathematical and Computational Foundations of Learning Theory
Matthias Hein (Universität des Saarlandes, DE), Gabor Lugosi (UPF – Barcelona, ES), Lorenzo Rosasco (MIT – Cambridge, US)
August 30 to September 4, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15361

15371 – Quantum Cryptanalysis
Michele Mosca (University of Waterloo, CA), Martin Roetteler (Microsoft Corporation – Redmond, US), Nicolas Sendrier (INRIA – Le Chesnay, FR), Rainer Steinwandt (Florida Atlantic University – Boca Raton, US)
September 6–11, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15371

15381 – Information from Deduction: Models and Proofs
Nikolaj S. Bjorner (Microsoft Corporation – Redmond, US), Jasmin Christian Blanchette (INRIA Lorraine – Nancy, FR), Viorica Sofronie-Stokkermans (Universität Koblenz-Landau, DE), Christoph Weidenbach (MPI für Informatik – Saarbrücken, DE)
September 13–18, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15381

15382 – Modeling and Simulation of Sport Games, Sport Movements, and Adaptations to Training
Ricardo Duarte (University of Lisbon, PT), Björn Eskofier (Universität Erlangen-Nürnberg, DE), Patrick Lucey (Disney Research – Pittsburgh, US), Martin Rumpf (Universität Bonn, DE), Josef Wiemeyer (TU Darmstadt, DE)
September 13–16, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15382

15391 – Algorithms and Complexity for Continuous Problems
Aicke Hinrichs (Universität Linz, AT), Joseph F. Traub (New York, US), Henryk Wozniakowski (Columbia University – New York, US), Larisa Yaroslavtseva (Universität Passau, DE)
September 20–25, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15391

15392 – Measuring the Complexity of Computational Content: Weihrauch Reducibility and Reverse Analysis
Vasco Brattka (Universität der Bundeswehr – München, DE), Akitoshi Kawamura (University of Tokyo, JP), Alberto Marcone (University of Udine, IT), Arno Pauly (University of Cambridge, GB)
September 20–25, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15392

15401 – Circuits, Logic and Games
Mikolaj Bojanczyk (University of Warsaw, PL), Meena Mahajan (The Institute of Mathematical Sciences, IN), Thomas Schwentick (TU Dortmund, DE), Heribert Vollmer (Leibniz Universität Hannover, DE)
September 27 to October 2, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15401

15402 – Self-Assembly and Self-Organization in Computer Science and Biology
Vincent Danos (University of Edinburgh, GB), Heinz Koeppl (TU Darmstadt, DE)
September 27 to October 2, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15402

15411 – Multimodal Manipulation Under Uncertainty
Jan Peters (TU Darmstadt, DE), Justus Piater (Universität Innsbruck, AT), Robert Platt (Northeastern University – Boston, US), Siddhartha Srinivasa (Carnegie Mellon University, US)
October 4–9, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15411

15412 – Dynamic Traffic Models in Transportation Science
José R. Correa (University of Chile – Santiago de Chile, CL), Tobias Harks (Maastricht University, NL), Kai Nagel (TU Berlin, DE), Britta Peis (RWTH Aachen, DE), Martin Skutella (TU Berlin, DE)
October 4–9, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15412

15421 – Rack-Scale Computing
Babak Falsafi (EPFL – Lausanne, CH), Tim Harris (Oracle Labs – Cambridge, GB), Dushyanth Narayanan (Microsoft Research UK – Cambridge, GB), David Patterson (University of California – Berkeley, US)
October 11–16, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15421

15431 – Genomic Privacy
Jean Pierre Hubaux (EPFL – Lausanne, CH), Stefan Katzenbeisser (TU Darmstadt, DE), Bradley Malin (Vanderbilt University – Nashville, US), Gene Tsudik (University of California – Irvine, US)
October 18–23, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15431

15441 – Duality in Computer Science
Mai Gehrke (University of Paris VII, FR), Achim Jung (University of Birmingham, GB), Victor Selivanov (A. P. Ershov Institute – Novosibirsk, RU), Dieter Spreen (Universität Siegen, DE)
October 25–30, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15441

15442 – Approaches and Applications of Inductive Programming
Jose Hernandez-Orallo (Technical University of Valencia, ES), Stephen H. Muggleton (Imperial College London, GB), Ute Schmid (Universität Bamberg, DE), Benjamin Zorn (Microsoft Research – Redmond, US)
October 25–30, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15442

15451 – Verification of Evolving Graph Structures
Parosh Aziz Abdulla (Uppsala University, SE), Fabio Gadducci (University of Pisa, IT), Barbara König (Universität Duisburg-Essen, DE), Viktor Vafeiadis (MPI-SWS – Kaiserslautern, DE)
November 1–6, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15451

15461 – Vision for Autonomous Vehicles and Probes
Andrés Bruhn (Universität Stuttgart, DE), Atsushi Imiya (Chiba University, JP), Ales Leonardis (University of Birmingham, GB), Tomas Pajdla (Czech Technical University, CZ)
November 8–13, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15461

15462 – The Mobile Revolution – Machine Intelligence for Autonomous Vehicles
Wolfram Burgard (Universität Freiburg, DE), Markus Enzweiler (Daimler AG – Böblingen, DE), Uwe Franke (Daimler AG – Sindelfingen, DE), Mohan Trivedi (University of California, San Diego – La Jolla, US)
November 10–13, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15462

15471 – Symbolic Computation and Satisfiability Checking
Erika Abraham (RWTH Aachen, DE), Pascal Fontaine (LORIA – Nancy, FR), Thomas Sturm (MPI für Informatik – Saarbrücken, DE), Dongming Wang (Beihang University – Beijing, CN)
November 15–20, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15471

15472 – Programming with "Big Code"
William W. Cohen (Carnegie Mellon University, US), Charles Sutton (University of Edinburgh, GB), Martin T. Vechev (ETH Zürich, CH)
November 15–18, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15472

15481 – Evaluation in the Crowd: Crowdsourcing and Human-Centred Experiments
Daniel Archambault (Swansea University, GB), Tobias Hoßfeld (Universität Duisburg-Essen, DE), Helen C. Purchase (University of Glasgow, GB)
November 22–27, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15481

15482 – Social Concepts in Self-Organising Systems
Ada Diaconescu (Telecom Paris Tech, FR), Stephen Marsh (UOIT – Oshawa, CA), Jeremy Pitt (Imperial College London, GB), Wolfgang Reif (Universität Augsburg, DE), Jan-Philipp Steghöfer (Chalmers UT – Göteborg, SE)
November 22–27, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15482

15491 – Approximate and Probabilistic Computing: Design, Coding, Verification
Antonio Filieri (Imperial College London, GB), Marta Kwiatkowska (University of Oxford, GB), Sasa Misailovic (MIT – Cambridge, US), Todd Mytkowicz (Microsoft Corporation – Redmond, US), Martin C. Rinard (MIT – Cambridge, US)
November 29 to December 4, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15491

15492 – Computational Metabolomics
Sebastian Böcker (Universität Jena, DE), Oliver Fiehn (University of California – Davis, US), Juho Rousu (Aalto University, FI), Emma Schymanski (Eawag – Dübendorf, CH)
November 29 to December 4, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15492

15511 – The Graph Isomorphism Problem
Laszlo Babai (University of Chicago, US), Anuj Dawar (University of Cambridge, GB), Pascal Schweitzer (RWTH Aachen, DE), Jacobo Torán (Universität Ulm, DE)
December 13–18, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15511

15512 – Debating Technologies
Iryna Gurevych (TU Darmstadt, DE), Eduard H. Hovy (Carnegie Mellon University – Pittsburgh, US), Noam Slonim (IBM – Haifa, IL), Benno Stein (Bauhaus-Universität Weimar, DE)
December 13–18, 2015 | Dagstuhl Seminar | http://www.dagstuhl.de/15512

Dagstuhl-Perspektiven-Workshops 15.2 *Dagstuhl Perspectives Workshops*

15302 – Digital Scholarship and Open Science in Psychology and the Behavioral Sciences
Alexander Garcia Castro (Technical University of Madrid, ES), Janna Hastings (European Bioinformatics Institute – Cambridge, GB), Robert Stevens (University of Manchester, GB), Erich Weichselgartner (ZPID – Trier, DE)
July 19–24, 2015 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/15302

15342 – Power-Bounded HPC Performance Optimization
Dieter Kranzlmüller (LMU München, DE), Barry L. Rountree (LLNL – Livermore, US)
August 16–21, 2015 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/15342

15362 – Present and Future of Formal Argumentation
Dov M. Gabbay (King's College London, GB), Massimiliano Giacomin (University of Brescia, IT), Beishui Liao (Zhejiang University, CN), Leon van der Torre (University of Luxembourg, LU)
August 30 to September 4, 2015 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/15362

15452 – Artifact Evaluation for Publications
Bruce R. Childers (University of Pittsburgh, US), Grigori Fursin (cTuning – Cachan, FR), Shriram Krishnamurthi (Brown University – Providence, US), Andreas Zeller (Universität des Saarlandes, DE)
November 1–4, 2015 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/15452

GI-Dagstuhl-Seminare 15.3 *GI-Dagstuhl Seminars*

15212 – Computational Imaging
Matthias B. Hullin (Universität Bonn, DE), Ivo Ihrke (INRIA – Bordeaux, FR), Lars Omlor (Carl Zeiss – Oberkochen, DE)
May 17–22, 2015 | GI-Dagstuhl Seminar | http://www.dagstuhl.de/15212

15283 – Entertainment Computing and Serious Games
Ralf Dörner (Hochschule RheinMain – Wiesbaden, DE), Stefan Göbel (TU Darmstadt, DE), Katharina A. Zweig (TU Kaiserslautern, DE)
July 5–10, 2015 | GI-Dagstuhl Seminar | http://www.dagstuhl.de/15283

15504 – Formal Evaluation of Critical Infrastructures
Erika Abraham (RWTH Aachen, DE), Anne Remke (Universität Münster, DE), Markus Siegle (Universität der Bundeswehr – München, DE), Marielle Stoelinga (University of Twente, NL)
December 6–9, 2015 | GI-Dagstuhl Seminar | http://www.dagstuhl.de/15504

Lehrveranstaltungen 15.4 *Educational Events*

15092 – Winter School "Database Implementation Techniques"
Goetz Graefe (HP Labs – Madison, US)
February 22–27, 2015 | Educational Event | http://www.dagstuhl.de/15092

15503 – Lehrerfortbildung in Informatik
Roswitha Bardohl (Schloss Dagstuhl – Saarbrücken, DE), Manuel Garcia Mateos (LPM Saarbrücken, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 9–11, 2015 | Educational Event | http://www.dagstuhl.de/15503

Forschungsgruppentreffen 15.5 *Research Group Meetings*

15044 – Erneuerbare Mobilität – Intermodales Reisen
Karl-Heinz Krempels (RWTH Aachen, DE), Christoph Terwelp (RWTH Aachen, DE)
January 21–23, 2015 | Research Group Meeting | http://www.dagstuhl.de/15044

15073 – Klausurtagung AG Hettel
Jörg Hettel (HS Kaiserslautern – Zweibrücken, DE)
February 11–13, 2015 | Research Group Meeting | http://www.dagstuhl.de/15073

15103 – Lehrstuhltreffen AG Zeller
Andreas Zeller (Universität des Saarlandes, DE)
March 4–6, 2015 | Research Group Meeting | http://www.dagstuhl.de/15103

15104 – Lehrstuhltreffen LST Zitterbart
Sören Finster (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)
March 4–6, 2015 | Research Group Meeting | http://www.dagstuhl.de/15104

15132 – Klausurtagung AG Hartenstein
Hannes Hartenstein (KIT – Karlsruher Institut für Technologie, DE)
March 25–27, 2015 | Research Group Meeting | http://www.dagstuhl.de/15132

15133 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)
Michaela Huhn (TU Clausthal, DE), Matthias Riebisch (Universität Hamburg, DE)
March 22–25, 2015 | Research Group Meeting | http://www.dagstuhl.de/15133

15135 – Klausurtagung AG Schneider
Klaus Schneider (TU Kaiserslautern, DE)
March 26–27, 2015 | Research Group Meeting | http://www.dagstuhl.de/15135

15143 – Klausurtagung LST Halang
Wolfgang A. Halang (FernUniversität in Hagen, DE)
March 29 to April 1, 2015 | Research Group Meeting | http://www.dagstuhl.de/15143

15144 – GIBU 2015: GI-Beirat der Universitätsprofessoren
Gregor Snelting (KIT – Karlsruher Institut für Technologie, DE)
March 29 to April 1, 2015 | Research Group Meeting | http://www.dagstuhl.de/15144

15149 – Forschungsaufenthalt
Stephan Diehl (Universität Trier, DE)
March 31 to April 1, 2015 | Research Group Meeting | http://www.dagstuhl.de/15149

15202 – Kolloquium zum GI Dissertationspreis 2014
Steffen Hölldobler (TU Dresden, DE)
May 10–13, 2015 | Research Group Meeting | http://www.dagstuhl.de/15202

15232 – Gemeinsamer Workshop der Graduiertenkollegs – 9th Joint Workshop of the German Research Training Groups in Computer Science: GRK 1773 and GRK 1763
Parvaneh Babari (Universität Leipzig, DE), Sergiu Dotenco (Universität Erlangen – Nürnberg, DE), Rainer Müller (Universität Erlangen-Nürnberg, DE), Marc Stamminger (Universität Erlangen-Nürnberg, DE), Anna Yupatova (Universität Erlangen-Nürnberg, DE)
May 31 to June 3, 2015 | Research Group Meeting | http://www.dagstuhl.de/15232

15253 – Nachfolge-Workshop Buchprojekt "Corporate Semantic Web"
Börteçin Ege (TU Wien, AT), Thomas Hoppe (Datenlabor Berlin, DE), Bernhard Humm (Hochschule Darmstadt, DE), Anatol Reibold (OntoPort UG – Darmstadt, DE)
June 14–17, 2015 | Research Group Meeting | http://www.dagstuhl.de/15253

15273 – Retreat SFB 1102: Information Density and Linguistic Encoding
Elke Teich (Universität des Saarlandes, DE)
June 28–30, 2015 | Research Group Meeting | http://www.dagstuhl.de/15273

15275 – Lehrstuhltreffen Rechtsinformatik
Christoph Sorge (Universität des Saarlandes, DE)
June 29–30, 2015 | Research Group Meeting | http://www.dagstuhl.de/15275

15276 – Computer-Literacy im Fachunterricht der Sekundar-Stufe I
Jens Bochmann (IGS / Realschule plus Georg-Friedrich-Kolb Speyer, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
June 29 to July 1, 2015 | Research Group Meeting | http://www.dagstuhl.de/15276

15292 – GI-Klausur zu "Grand Challenges" der Informatik
Simone Rehm (Trumpf – Ditzingen, DE), Cornelia Winter (GI Bonn, DE)
July 12–15, 2015 | Research Group Meeting | http://www.dagstuhl.de/15292

15293 – Lehrstuhltreffen LST Sick
Bernhard Sick (Universität Kassel, DE)
July 12–16, 2015 | Research Group Meeting | http://www.dagstuhl.de/15293

15294 – The Continuing Arms Race: Code-Reuse Attacks and Defenses
Ulfar Erlingsson (Google Inc. – Mountain View, US), Thorsten Holz (Ruhr-Universität Bochum, DE), Per Larsen (University of California – Irvine, US), Ahmad-Reza Sadeghi (TU Darmstadt, DE)
July 12–15, 2015 | Research Group Meeting | http://www.dagstuhl.de/15294

15295 – Scalable Author Disambiguation for Bibliographic Databases: Project Kick-Off
Marcel R. Ackermann (Schloss Dagstuhl – Wadern, DE)
July 15–16, 2015 | Research Group Meeting | http://www.dagstuhl.de/15295

15313 – Klausurtagung LST Freiling
Felix C. Freiling (Universität Erlangen-Nürnberg, DE)
July 26–29, 2015 | Research Group Meeting | http://www.dagstuhl.de/15313

15314 – Deutsch-Pakistanischer Workshop
Karsten Berns (TU Kaiserslautern, DE)
July 30–31, 2015 | Research Group Meeting | http://www.dagstuhl.de/15314

15315 – Fakultätsklausur der Fakultät für Informatik – SRH Hochschule Heidelberg
Barbara Sprick (SRH Hochschule Heidelberg, DE)
July 27–28, 2015 | Research Group Meeting | http://www.dagstuhl.de/15315

15423 – DDI: Facilitating Process and Metadata-Driven Automation in the Social, Economic, and Behavioural Sciences with the Data Documentation Initiative
Michelle Edwards (Cornell University, US), Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (University of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 11–16, 2015 | Research Group Meeting | http://www.dagstuhl.de/15423

15433 – DDI Moving Forward: Facilitating Interoperability and Collaboration with Other Metadata Standards
Arofan Gregory (Open Data Foundation – Tucson, US), Wendy Thomas (University of Minnesota – Minneapolis, US), Mary Vardigan (University of Michigan – Ann Arbor, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 18–23, 2015 | Research Group Meeting | http://www.dagstuhl.de/15433

15439 – Forschungsaufenthalt: Regelbasierte Systeme
Karl-Heinz Krempels (RWTH Aachen, DE)
October 19–21, 2015 | Research Group Meeting | http://www.dagstuhl.de/15439

15453 – Buchprojekt: Modellierung in der Informatik
Bernhard Rumpe (RWTH Aachen, DE), Bernhard Thalheim (Universität Kiel, DE)
November 4–6, 2015 | Research Group Meeting | http://www.dagstuhl.de/15453

15473 – Klausurtagung LST Schmeck
Birger Becker (FZI – Karlsruhe, DE), Hartmut Schmeck (KIT – Karlsruher Institut für Technologie, DE)
November 18–20, 2015 | Research Group Meeting | http://www.dagstuhl.de/15473

15505 – Secan Lab Seminar
Thomas Engel (University of Luxembourg, LU)
December 7–8, 2015 | Research Group Meeting | http://www.dagstuhl.de/15505

15509 – Forschungsaufenthalt: Regelbasierte Systeme
Karl-Heinz Krempels (RWTH Aachen, DE)
December 9–11, 2015 | Research Group Meeting | http://www.dagstuhl.de/15509

Mitglied der Leibniz-Gemeinschaft

Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2015 | ISSN 2199-1995
http://www.dagstuhl.de