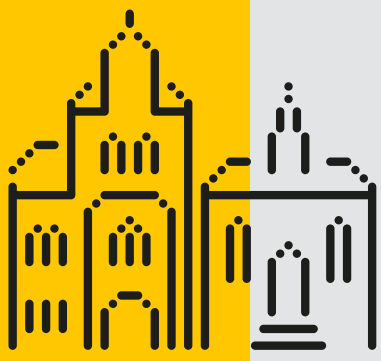

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2014

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Jahresbericht
Annual Report
2014

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. (Wissenschaftlicher Direktor) Heike Meißner (Technisch-administrative Geschäftsführerin)
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, Laura Cunniff, 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, Antónia Koštová, 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	Bonifatius GmbH, Paderborn
Druckauflage	160
ISSN	2199-1995
Online-Version	http://www.dagstuhl.de/dagpub/2199-1995
Veröffentlichungsdatum	Mai 2015

Leibniz-Gemeinschaft

Zertifikat seit 2012 audit berufundfamilie

Vorwort
Foreword

Dies ist mein erstes Vorwort für einen Dagstuhl Jahresbericht. Bis jetzt hat der Gründungsdirektor, Reinhard Wilhelm, solche Prolegomena geschrieben. Ende April des Berichtsjahres nach über 24 Jahren Aufbauarbeit hat er das Amt des Wissenschaftlichen Direktors an mich übergeben. Zwei Monate später kam es auch zu einem schon lange geplanten Wechsel in der technisch-administrativen Geschäftsführung, und Frau Heike Meißner hat Herrn Christian Lindig abgelöst.

Die neue Geschäftsführung hätte eigentlich keine schönere Aufgabe finden können. Sie fand eine äußerst wohl bestellte Institution vor: die Infrastruktur in hervorragendem Zustand, die Teams gut eingespielt, die Verfahren durch lange Erfahrung optimiert und wohl eingeübt. Und dazu gab es noch den für den Anfang so wichtigen Luxus, bei Unklarheiten die alte Geschäftsführung konsultieren zu können und bereitwillig Rat zu bekommen.

Das „Hauptgeschäft", die Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops, brummt. Aus den folgenden Seiten kann man sich einen guten Eindruck von der Vielfalt dieser Veranstaltungen machen. Auch in den anderen Aufgabenbereichen, die sich Dagstuhl gesetzt hat, ist im Berichtsjahr einiges Bemerkenswertes geschehen. Bei Dagstuhls Open-Access-Publishing Aktivitäten ist es bei der Tagungsbandserie LIPIcs zu einem Quantensprung gekommen. Die Zahl der betroffenen Tagungsreihen ist von 7 auf fast 20 gestiegen, mit so etablierten Konferenzen wie ICALP und ESA unter den Neuzugängen. Der in Zusammenarbeit mit der Universität Trier betriebenen Informatik-Publikationsdatenbank dblp wurde eine Art Ritterschlag verliehen. Beim deutschen Hochschulranking durch CHE und ZEIT für den Bereich Informatik werden die dblp Daten als eine der Hauptgrundlagen zur Bewertung der Forschungsleistungen herangezogen.

Das Jahr 2014 war für Dagstuhl ein Jahr des Übergangs, aber auch der Kontinuität. Die hervorragende Qualität von

This is my first preface for a Dagstuhl Annual Report. Previously, such prolegomena were written by Reinhard Wilhelm, the founding Director of Schloss Dagstuhl who – after 24 years of committed leadership and development– handed the position of Scientific Director over to me on May 1, 2014. Two months later, in a long-planned administrative change of hands, Heike Meißner succeeded Dr. Christian Lindig as Technical Administrative Director of Schloss Dagstuhl.

The new management could not have been given a more appealing job. We had inherited an exceptionally well-run institution: the infrastructure was in excellent condition, the staff team seasoned and highly trained, their work routines optimized through long experience. We were also given the luxury of expert consultancy with the center's former directors, who readily placed themselves at our disposition.

The center's main business, the Dagstuhl Seminar and Dagstuhl Perspectives Workshops, is humming. Readers will get a good impression of the diversity of these meetings on the following pages. In addition, other tasks areas that form part of Schloss Dagstuhl's core mission witnessed some remarkable developments during the year under review. The LIPIcs proceedings series, part of Dagstuhl's Open Access publishing platform, made a quantum leap forward in 2014 as the number of conferences supported by the series rose from 7 to nearly 20, with established conferences such as ICALP and ESA included among the newcomers. Schloss Dagstuhl's cooperative work with the dblp computer science database, operated from the University of Trier, also received an honor of sorts in 2014 when the German think tank CHE (Center for Higher Education Development) requested from dblp permission to use its data as one of the main bases for evaluating the research performance of scientists working in the field of computer science for a nationwide German university

Dagstuhl zu bewahren und in den neueren Aufgabenbereiche ähnliches Ansehen zu erlangen wie im Seminarbetrieb, wird eine interessante Herausforderung werden, der die neue Geschäftsführung mit Freude entgegensieht.

Raimund Seidel

Im Namen der Geschäftsführung

Prof. Raimund Seidel, Ph. D.
Wissenschaftlicher Direktor

Heike Meißner
Technisch-administrative Geschäftsführerin

ranking. The ranking will be published in the German news weekly, Die Zeit, foreseeably in 2015.

2014 was a year of transition for Dagstuhl, but also of continuity. To maintain the excellent quality of Dagstuhl's scientific program and match that in newer areas of responsibility is an interesting challenge that the new management looks forward to with joy.

On behalf of the Managing Directors

Inhaltsverzeichnis
Contents

1 Das Zentrum Schloss Dagstuhl
Schloss Dagstuhl Center

Dagstuhl's Mission 1.1 *Dagstuhl's Mission*

Dagstuhl Seminar Program

Schloss Dagstuhl, Leibniz-Zentrum für Informatik, hat zum Ziel, Informatikforschung von international anerkanntem Rang zu fördern, Weiterbildung auf hohem fachlichen Niveau durchzuführen und den Wissensaustausch zwischen Forschung und Praxis anzuregen. Das Zentrum veranstaltet hierzu Forschungsseminare, die führende Wissenschaftler aus der ganzen Welt für eine Woche zu einem intensiven Wissensaustausch zusammenführen. Die Seminare ermöglichen den Austausch und die Entwicklung neuer Ideen, die Diskussion aktueller Probleme sowie die Weichenstellung für zukünftige Entwicklungen.

Die Idee zur Gründung von Schloss Dagstuhl 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.

Das wissenschaftliche Programm von Schloss Dagstuhl umfasst die sogenannten Dagstuhl-Seminare sowie Dagstuhl-Perspektiven-Workshops. Dem vielversprechenden Forschungsnachwuchs wird dabei die Möglichkeit gegeben, in seinem speziellen Fachgebiet mit exzellenten Experten zusammenzuarbeiten und neue Sichtweisen zu diskutieren. Das Programm eines Dagstuhl-Seminars wird absichtlich flexibel gestaltet, um eine gemeinschaftliche Atmosphäre zu schaffen, die in dynamischer Weise offene und kreative Diskussionen zulässt. Bei einem Dagstuhl-Perspektiven-Workshop hingegen diskutiert eine oftmals kleinere Gruppe von ausgewiesenen Experten ein Themengebiet und seine perspektivische Ausrichtung. Hierzu wird die aktuelle Situation eines Forschungsgebietes analysiert, um darauf aufbauend strategische Empfehlungen und richtungsweisende Perspektiven für die weitere Zukunft zu entwickeln. Die Erkenntnisse werden in einem Manifest zusammengefasst, das auch an (politische) Entscheidungsträger weitergegeben wird.

Die Seminare und Perspektiven-Workshops werden jeweils von einer kleinen Gruppe ausgewiesener Wissenschaftler im entsprechenden Gebiet beantragt. Für die Begutachtung der Vorschläge und der Teilnehmerlisten ist das Wissenschaftliche Direktorium verantwortlich, bevor Anträge akzeptiert und in Dagstuhls wissenschaftliches Programm aufgenommen werden. Die Teilnahme an diesen Veranstaltungen ist nur mit einer persönlichen Einladung durch das Zentrum möglich. Um den besten internationalen Wissenschaftlern eine Teilnahme zu ermöglichen, wird ein Teil der Aufenthaltskosten von Dagstuhl übernommen.

Schloss Dagstuhl befindet sich in einer ländlichen Gegend im nördlichen Saarland, im Herzen des Drei-

Dagstuhl Seminar Program

The mission of the Leibniz-Zentrum für Informatik at Schloss Dagstuhl is to promote world-class research in informatics, support cutting-edge continuing education and professional development, and encourage the exchange of knowledge and findings between the academic community and industry. The center hosts research seminars in which leading researchers from all over the world live and work together at Schloss Dagstuhl for several days in an intensive research climate. New ideas are showcased, topical problems are discussed, and the course is set for future development in the field.

The idea behind Schloss Dagstuhl 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. Dagstuhl was founded in 1990 and quickly became established as one of the world's premier centers for informatics research.

The center's scientific program includes the well-known Dagstuhl Seminars series and its complement, the Dagstuhl Perspectives Workshops series. Dagstuhl Seminars offer promising young researchers in a specific cutting-edge field of informatics the opportunity to work closely together and share their views and findings with the international elite of their field. The seminars thrive on an open-ended program that allows participants to take advantage of synergies as they come up over several days, creating a dynamic space for discussion and debate that often leads in unexpected directions. By contrast, Dagstuhl Perspectives Workshops bring together a group of well-established senior researchers to discuss a topic area and its perspectives. The goal is to analyze the overall state of the field in order to detect strategic trends and develop new perspectives on its continued evolution. The results are collected and published in a Dagstuhl Manifesto, which is made available to policymakers.

Each Dagstuhl Seminar and Dagstuhl Perspectives Workshop is headed by a small group of scientists of international standing in their respective fields. Proposals for seminars and workshops are reviewed by the Dagstuhl Scientific Directorate before their acceptance into the center's scientific program. Participation in these events is possible by way of personal invitation only by the center, which assumes part of the associated costs in order to enable the world's most qualified scientists to participate.

Located 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

länderecks 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 der besten Informatik-Fachbibliotheken weltweit, sowie eine Vielzahl von zusätzlichen Arbeits- und Freizeiträumen bieten vielfältige Möglichkeiten, damit sich die Gäste auch außerhalb des fachlichen Seminarprogramms kennenlernen und austauschen können.

Dagstuhl Publishing

Es gehört zur zentralen Aufgabe von Schloss Dagstuhl, die Kommunikation zwischen den Wissenschaftlern in der Informatik zu fördern. 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 haben.

Bibliographiedatenbank dblp

Seit 2011 betreibt Schloss Dagstuhl in enger Zusammenarbeit mit der Universität Trier die Literaturdatenbank dblp, welche mit etwa 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 verfügbar machen von personalisierten Publikationsnachweisen. Durch den dblp-Dienst leistet Schloss Dagstuhl einen weiteren Beitrag im Rahmen seiner Mission zur Förderung der Erkennung, Verbreitung und Umsetzung neuer Informatikerkenntnisse auf international anerkanntem Niveau.

Finanzierung

Schloss Dagstuhl wird durch eine Bund-Länder-Förderung finanziert und beherbergt jedes Jahr mehr als 3 500 internationale Gäste. Seit 2006 ist Schloss Dagstuhl Mitglied in der Leibniz-Gemeinschaft, einem Verbund von 89 Forschungsinstituten, Bibliotheken und Museen.[1]

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.

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.

dblp computer science bibliography

In parallel to its seminar program, Schloss Dagstuhl has operated the dblp computer science bibliography together with the University of Trier since 2011. 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.

Funding

Schloss Dagstuhl is jointly funded by the German federal and state governments and hosts over 3,500 research guests each year from countries across the globe. Since 2006, it has been a member of the Leibniz Association, a non-profit research consortium composed of 89 research institutes, libraries and museums throughout Germany.[1]

[1] Stand: Dezember, 2014
As of December, 2014

Fig. 1.1
Schloss Dagstuhl's historical „Schloss" is the soul of the center.

Fig. 1.2
The center received the Bauherrenpreis architectural award in 2014 for its modern guest house, shown here in the foreground.

Neuigkeiten in 2014 1.2 *News from 2014*

Das Team

Das Jahr 2014 brachte tiefgreifende Veränderungen für das Dagstuhl-Team mit sich. Am 30. April ging Gründungsdirektor Prof. Reinhard Wilhelm nach fast 25 Jahren als Wissenschaftlicher Direktor in den Ruhestand. Von der Gründung im Jahr 1990 an hat er das Zentrum zu seiner heutigen herausragenden Bedeutung geführt. Dabei spielte er eine große Rolle bei der Entwicklung und Abstimmung des „Dagstuhl-Konzepts“, das den Schwerpunkt auf einen (ergebnis-)offenen Wissens- und Ideenaustausch in entspannter Atmosphäre legt. Prof. Wilhelms Weitblick hatte großen Einfluss auf die Leitung des Zentrums und ermöglichte Schloss Dagstuhl während der 1990er und frühen 2000er Jahre, als sich die Informatik als wissenschaftliche Disziplin extrem schnell weiterentwickelte, die Beibehaltung seiner ursprünglichen Bestimmung.

Mit dem Ausscheiden Prof. Wilhelms wurde die Leitung des Zentrums an Prof. Raimund Seidel übergeben, der am 1. Mai 2014 das Amt des Wissenschaftlichen Direktors antrat. Raimund Seidel, ein international renommierter Wissenschaftler, Professor für Informatik und Gründungssprecher der „Saarbrücken Graduate School of Computer Science“ an der Universität des Saarlandes, studierte Mathematik und Informatik in Österreich, Kanada und den USA, wo er seinen Doktortitel von der Cornell University erhielt, und lehrte von 1987 bis 1994 an der University of California in Berkeley. Seine Interessen liegen hauptsächlich im Entwurf und der Analyse von Algorithmen und Datenstrukturen, insbesondere für geometrische Probleme und auch unter Verwendung von Zufallsmechanismen („Randomisierung“). Abgesehen von seiner langjährigen Tätigkeit als Mitglied des Wissenschaftlichen Direktoriums von Schloss Dagstuhl und seiner regelmäßigen Teilnahme am wissenschafltichen Programm des Zentrums ist Prof. Seidel maßgeblich an den Erfolgen der Saarbrücker Informatik in der Exzellenzinitiative des Bundesministeriums für Bildung und Forschung und der Deutschen Forschungsgemeinschaft beteiligt.

Am 18. Juli 2014 fanden sich Wissenschaftler, Mitarbeiter und enge Freunde auf Schloss Dagstuhl zusammen, um Reinhard Wilhelms herausragende Leistung zu würdigen und Raimund Seidel als seinen Nachfolger willkommen zu heißen. Es fand ein Festkolloquium mit eingeladenen Vorträgen von Prof. Moshe Y. Vardi (Rice University), Prof. Friedrich Eisenbrand (École Polytechnique Fédérale de Lausanne), Prof. Andreas Reuter (Heidelberger Institut für Theoretische Studien) und Prof. Otto Spaniol (RWTH Aachen) statt, gefolgt von einem abendlichen Bankett im Schlossgarten.

Dieser Führungswechsel war auch Thema mehrerer Berichte in den deutschen Medien. In einem Interview äußerte sich Raimund Seidel optimistisch zur Zukunft von Schloss Dagstuhl und zu den bevorstehenden Herausforderungen. Der globale Trend zur computervermittelten Kom-

The Team

The year 2014 was a profoundly transitional one for the Dagstuhl team. On April 30 Schloss Dagstuhl's founding director, Reinhard Wilhelm, retired from the position of Scientific Director after nearly 25 years of leadership. Prof. Wilhelm, who built the center up from scratch in 1990 and steered it into a position of preeminence during the ensuing decades, is largely responsible for developing and fine-tuning the "Dagstuhl concept," with its emphasis on the open-ended exchange of scientific knowledge and ideas in a naturally relaxed setting. His vision has deeply influenced the direction of the center and enabled Schloss Dagstuhl to stay true to its original purpose throughout the 1990s and early 2000s, when computer science itself was rapidly evolving as a discipline.

With the retirement of Prof. Wilhelm, the directorship of the center was passed on to Raimund Seidel, who acceded as Scientific Director of Schloss Dagstuhl on May 1, 2014. An internationally renowned scientist, professor of informatics and founding chair of the Graduate School of Computer Science at Saarland University, Prof. Seidel studied mathematics and computer science in Austria, Canada, and the United States, where he received his Ph. D. from Cornell University and taught at the University of California at Berkeley from 1987 to 1994. Prof. Seidel's research interests focus mainly on modeling and analyzing algorithms and data structures, especially in the context of geometric problems, using randomization. In addition to his longstanding support of Schloss Dagstuhl as a former member of its Scientific Directorate and a frequent participant in its scientific program, Prof. Seidel has made significant contributions to the success of informatics research in Saarbrücken within the context of the Excellence Initiative of the German Federal Ministry of Education and Research and the German Research Foundation.

On July 18, 2014, scientists, collaborators and close friends gathered together at Schloss Dagstuhl to honor the great work of Reinhard Wilhelm and to welcome Raimund Seidel as his successor. The event included a scientific program with invited talks from Prof. Moshe Y. Vardi from Rice University, Prof. Friedrich Eisenbrand from the École Polytechnique Fédérale de Lausanne, Prof. Andreas Reuter from the Heidelberger Institut für Theoretische Studien, and Prof. Otto Spaniol from RWTH Aachen, followed by an evening banquet in the Schloss garden.

The shift into new leadership was also the subject of several targeted reports in the German media. Commenting on the future of Schloss Dagstuhl and the challenges that lie ahead, Raimund Seidel noted optimistically in one interview that the global trend towards computer-mediated communication does not seem to have dampened scientists' enthusiasm for Schloss Dagstuhl with its focus on personal, face-to-face exchange.

munikation habe die Begeisterung der Wissenschaftler für Schloss Dagstuhl und die Ausrichtung auf persönlichen, direkten Austausch anscheinend nicht gedämpft.

Das Jahr 2014 brachte auch das schon lange geplante Ende der Amtszeit Dr. Christian Lindigs als Technisch-Administrativer Geschäftsführer am 30. Juni 2014. Die Stelle wurde am 1. Juli 2014 mit Heike Meißner neu besetzt. Dr. Lindig kam 2007 als Mitglied des wissenschaftlichen Stabs zum Dagstuhl-Team hinzu und wurde 2009 einer der Geschäftsführer der Gesellschaft. Seine Nachfolgerin Frau Meißner ist Diplom-Betriebswirtin (FH) und verfügt durch ihre Tätigkeiten bei verschiedenen deutschen und internationalen Unternehmen über umfangreiche Berufserfahrung im administrativen Bereich.

Mit diesen Ernennungen endeten die Veränderungen im Team, die 2012 mit der Pensionierung von Angelika Mueller-von Brochowski (1991–2012) und Wolfgang Lorenz (1990–2012) begonnen hatten. Frau Mueller und Herr Lorenz waren seit der Gründung des Zentrums für administrative Angelegenheiten zuständig.

Alle Mitarbeiter von Schloss Dagstuhl wurden 2014 über den Kernhaushalt des Zentrums bezahlt. Eine Ausnahme bildet das dblp-Team, das teilweise durch eine großzügige Spende der Klaus Tschira Stifung in Höhe von 67 800 € unterstützt wurde. Ende 2014 beschäftigte Schloss Dagstuhl insgesamt 31,3 Vollzeitäquivalente bzw. 44 Angestellte.

Dagstuhl Publishing

Die Open-Access-Publikationsaktiväten haben in 2014 starken Zuspruch bekommen, insbesondere in der Konferenzbandreihe LIPIcs, in welcher mehrere hochrangige wissenschaftliche Konferenzen, darunter z.B. ESA, ICALP und SoCG, aufgenommen wurden.

Mehr Informationen zu den Open-Access-Aktivitäten von Schloss Dagstuhl finden sich in Kapitel 6.

Bibliographiedatenbank dblp

Bereits seit 2011 besteht die Kooperation zwischen Schloss Dagstuhl und der Universität Trier für den gemeinsamen Betrieb von dblp. Ziel war und ist es, die thematische Breite der Datenbank zu vergrößern und das Angebot langfristig zu verstetigen und auszubauen. Dabei legt das dblp-Team großen Wert auf Verlässlichkeit und Qualität der Einträge.

Der Erfolg der Zusammenarbeit kann sich auch in Zahlen sehen lassen: Jeden Monat verzeichnet die dblp-Webseite bis zu 10 Millionen Zugriffe von über 450 000 verschiedenen Nutzern aus aller Welt. Dies entspricht etwa drei Seitenzugriffen pro Sekunde, und etwa alle drei Sekunden beginnt eine neue Nutzer-Session. Der Datenbestand wuchs dabei im Jahr 2014 um mehr als 350 000 Einträge. Ende 2014 indexierte dblp somit bereits über 2,8 Millionen Fachartikel aus den verschiedenen Teilgebieten der Informatik.

Die Datensauberkeit und -qualität stetig zu steigern ist eines der Hauptanliegen des dblp-Teams. Daher ist das

The year 2014 also marked the long-planned end of Dr. Christian Lindig's 5-year tenure as Technical Administrative Director on June 30, 2014. Dr. Lindig, who had joined Schloss Dagstuhl as a member of its scientific staff in 2007 before becoming a director of the company in 2009, left behind a record of committed service and leadership. He was succeeded by Heike Meißner on July 1, 2014. Ms. Meißner holds a degree in Business Administration and has extensive experience in that field with various national and international companies.

With these appointments, the team transition that had begun in 2012 with the retirement of Angelika Mueller-von Brochowski (1991–2012) and Wolfgang Lorenz (1990–2012), two of the center's founding figures on the administrative side, came to an end.

All staff at Schloss Dagstuhl were funded from the center's core budget in 2014, with the exception of the dblp team, which was partially supported in 2014 by a generous donation of 67,800 € from the Klaus Tschira Foundation. At the end of 2014, Schloss Dagstuhl had a total of 44 staff members including 31.3 full-time equivalent staff members.

Dagstuhl Publishing

Schloss Dagstuhl's open-access publishing services experienced a strong increase in demand from the community in 2014. This was especially true of the LIPIcs conference proceedings series, which received and accepted proposals from major scientific conferences such as ESA, ICALP and SoCG.

More information about the Open Access activities of Schloss Dagstuhl can be found in Chapter 6.

dblp computer science bibliography

In 2011, Schloss Dagstuhl and the University of Trier agreed to operate and maintain dblp jointly. The goal of the collaboration is to improve and extend dblp, and to guarantee the dblp service to the international computer science community long-term. One major focus of the dblp team is the reliability and quality of the provided data.

The success of the collaboration is already quite impressive: Up to ten million web pages are visited each month from more than 450,000 researchers all over the world. Every second, about three web pages are requested from the dblp web servers, and about every three seconds, a new user session is started. In 2014, the database grew by more than 350,000 new entries. By the end of 2014, dblp had already indexed about 2.8 million scholarly publications from all disciplines of computer science.

Improving reliability and quality of the provided data is always a main focus of the dblp team. Hence, the team is continuously improving and expanding its algorithmic tools. To improve the correct attribution of publications

Fig. 1.3
Outgoing and incoming directors at Schloss Dagstuhl on July 18, 2014. From left to right: Prof. Reinhard Wilhelm, Dr. Christian Lindig, Heike Meißner, and Prof. Raimund Seidel | Photo © Raphael Reischuk.

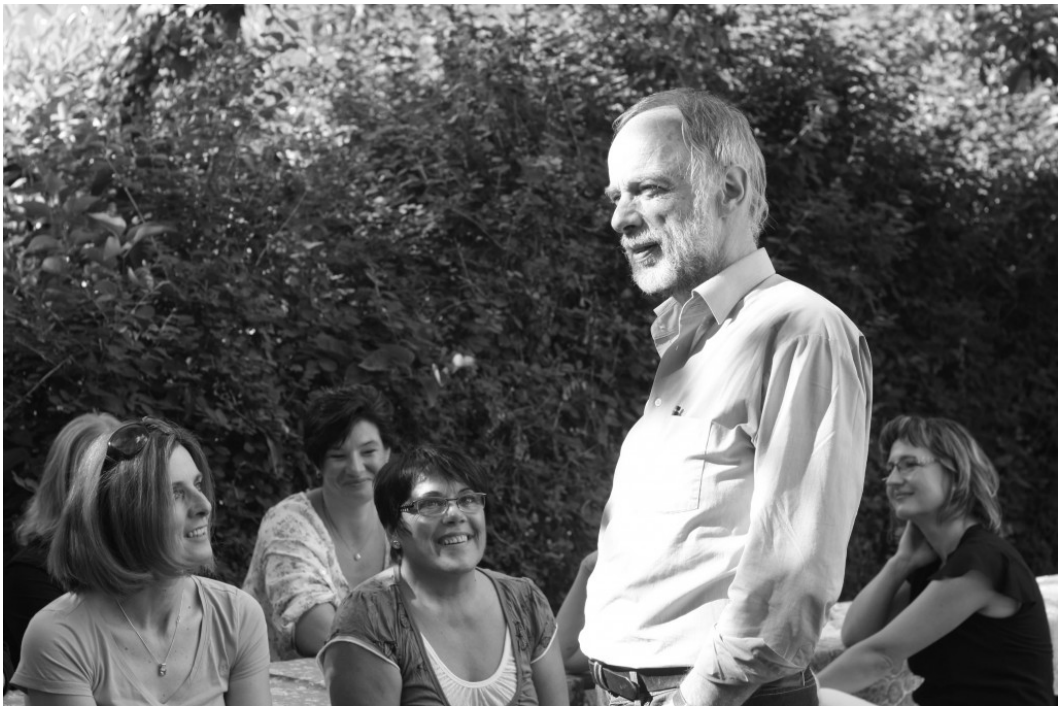

Fig. 1.4
Outgoing Scientific Director Prof. Reinhard Wilhelm, with staff in the summer of 2014. | Photo © Marc Herbstritt.

Team kontinuierlich dabei, seine algorithmischen Hilfsmittel zu verbessern und zu erweitern. Um die Urheberschaft wissenschaftlicher Publikationen eindeutig zu erkennen und zuzuordnen wurde 2014 eine Kooperation zwischen dblp, der mathematischen Literaturdatenbank zbMATH und dem Heidelberger Institut für Theoretische Studien (HITS) auf den Weg gebracht. Das Vorhaben konnte sich im „Leibniz Wettbewerb 2015“ durchsetzen und wird von der Leibniz Gemeinschaft für die kommenden drei Jahre gefördert.

dblp hat sich über die Jahre als die Referenz-Datenbank für bibliographische Informationen in der Informatik etabliert. Dem hat auch das Gütersloher Institut „Centrum für Hochschulentwicklung“ (CHE) Rechnung getragen und nutzt nun seit Ende 2014 dblp-Daten als Grundlage für bibliometrische Analysen innerhalb der Informatik im Rahmen seines „CHE Hochschulranking“.

Am 24. April 2014 hatten Dr. Michael Ley und Dr. Marcel R. Ackermann die Gelegenheit, Ministerin Doris Ahnen im Rahmen ihres Besuches an der Universität Trier die Zusammenarbeit zwischen der Universität und Schloss Dagstuhl näher vorzustellen. Dr. Ley präsentierte die Zusammenarbeit für die Literaturdatenbank dblp, zeigte die Vorteile der gemeinsamen Anstrengungen auf und verwies auf die Bestrebungen zur Etablierung einer langfristigen Zusammenarbeit.

Mehr Informationen zu dblp finden sich in Kapitel 7.

Öffentlichkeitsarbeit und Weiterbildung

Um junge Journalisten und Volontäre zu ermutigen, über anspruchsvolle Informatikthemen zu berichten, bietet Schloss Dagstuhl jährlich einen Workshop zum Thema Wissenschaftsjournalismus an. 2014 fand dieser parallel zu dem Dagstuhl-Seminar 14221 „Geometric Modeling“ vom 25. bis 28. Mai statt. Als Dozenten für den Workshop konnten wie in 2013 Tim Schröder (Wissenschaftsjournalist und Medientrainer, Oldenburg) und Gordon Bolduan (Pressesprecher des Exzellenz-Cluster „Multimodal Computing and Interaction“ an der Universität des Saarlandes) gewonnen werden. Alle Teilnehmer als auch die Dozenten waren höchst zufrieden mit den Inhalten und Ergebnissen des Workshops.

Schloss Dagstuhl engagiert sich im schulischen Bereich durch die Organisation einer jährlichen Lehrerfortbildung, die sich an Informatiklehrer 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. Das Interesse an dieser Fortbildung stieg seit dem Beginn in 1991 stetig an und die 24. „Lehrerfortbildung in Informatik“, die vom 10. bis 12. Dezember 2014 stattfand, führte mehr Teilnehmer zusammen als jemals zuvor. Die intensive Fortbildung richtet sich zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch häufen sich Anfragen zur Teilnahme von Lehrern aus anderen Bundesländern.

Zusätzlich zu diesen Weiterbildungsmöglichkeiten und Schulungen beteiligte sich das Zentrum an der interak-

to their unambiguous authors, a cooperation between dblp, the zbMATH database for mathematical literature, and the Heidelberg Institute for Theoretical Studies (HITS) has been initiated in 2014. The project has been successful in the “Leibniz Competition 2015” and will be funded by the Leibniz Association for three years.

Over the years, dblp has become the central resource of bibliographic meta data in computer science. This has also been acknowledged by the Gütersloh “Centre for Higher Education” (CHE) who started in 2014 using dblp data as basis for bibliometric analyses in computer science as part of its “CHE University Ranking”.

Schloss Dagstuhl’s support for the work of dblp through its strategic alliance with the University of Trier received recognition on April 24, 2014, when Dr. Michael Ley and Dr. Marcel R. Ackermann presented the cooperation to Minister Doris Ahnen during her visit of the University of Trier. Commenting on the future of the partnership, Dr. Ley pointed out that both partners view their teamwork as mutually advantageous and aim to achieve a stable long-term cooperation.

More information about dblp can be found in Chapter 7.

Public Relations and Professional Training

In order to encourage young journalists and trainees to report on complex informatics topics, Schloss Dagstuhl offers an annual workshop on science journalism. In 2014, the workshop took place on May 25–28 in parallel to Dagstuhl Seminar 14221, “Geometric Modeling.” Trainers included Tim Schröder from Oldenburg (scientific writer and media trainer) and Gordon Bolduan (press relations officer at the “Multimodal Computing and Interacting” Cluster of Excellence at the Saarland University). Participants as well as trainers and referees were very satisfied with the workshop.

Schloss Dagstuhl also offers a teacher training workshop specifically designed for teachers of 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). Interest in the workshop has risen steadily since the program began in 1991 and the 24th annual Dagstuhl teacher training workshop, held at Schloss Dagstuhl on December 10–12, 2014, attracted more participants than ever before. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, Schloss Dagstuhl does receive requests for participation from teachers of other federal states.

In addition to these professional training opportunities, the center also participated in the interactive exhibit “Begreifen und Verstehen” (“Grasp and Understand”) from September 25 to October 24 as part of the federally-spon-

tiven Ausstellung „Begreifen und Verstehen“, die vom 25. September bis zum 24. Oktober im Rahmen der durch den Bund geförderten Initiative „Wissenschaftsjahr 2014 – Die digitale Gesellschaft“ stattfand. Daneben war Schloss Dagstuhl vom 26. bis zum 28. November Teilnehmer bei der jährlichen Konferenz der Leibniz-Gesellschaft in Berlin vertreten. Bei beiden Veranstaltungen wurden Informationen zum wissenschaftlichen Auftrag und Programm des Zentrums bereitgestellt.

Mehr Informationen zur Öffentlichkeitsarbeit und zu den Weiterbildungsaktivitäten finden sich in Kapitel 5.

Zusammenarbeit mit dem Heidelberg Laureate Forum

2014 hatte Schloss Dagstuhl erneut die Möglichkeit zur Zusammenarbeit mit dem Heidelberg Laureate Forum[2] (HLF). Diese Veranstaltung bringt herausragende Mathematiker und Informatiker, nämlichGewinner des Abelpreises, der Fields-Medaille oder des ACM Turing Award, mit außergewöhnlich begabten jungen Wissenschaftlern aus aller Welt zusammen. Ausgewählte Teilnehmer des HLF 2014 – junge Wissenschaftler, die am Anfang ihrer Karriere stehen – erhielten im Rahmen der Zusammenarbeit die Gelegenheit zur Teilnahme an den Dagstuhl-Seminaren in der Woche vor der zweiten Ausgabe dieses Forums im September 2014. Es handelte sich dabei um die Dagstuhl-Seminare „Neural-Symbolic Learning and Reasoning“ (14381) und „Control Theory meets Software Engineering“ (14382) miteinbezogen wurden. Aufgrund des großen Erfolgs der Initiative haben alle Partner einer Fortsetzung der Zusammenarbeit für das Jahr 2015 zugestimmt.

In Vorbereitung auf diese gemeinsam getragene Woche begrüßte Schloss Dagstuhl am 13. Mai 2014 Dr. Klaus Tschira und einige Mitarbeiter des Heidelberger Teams zu einem außerordentlichen Besuch. Der gemeinschaftliche Austausch, ständig begleitet von Dr. Tschiras wohlüberlegten und kollegialen Anmerkungen, beinhaltete eine Führung durch das Konferenzzentrum, ein gemeinsames Mittagessen und Diskussionen am runden Tisch mit dem Wissenschaftlichen Direktor Prof. Raimund Seidel, dem ehemaligen Wissenschaftlichen Direktor Prof. Reinhard Wilhelm und einigen Mitgliedern der Stammbelegschaft von Schloss Dagstuhl. Das Treffen war Dr. Tschiras letzter offizieller Besuch auf Schloss Dagstuhl. Er verstarb unerwartet am 31. März 2015, als dieser Bericht gerade verfasst wurde. Schloss Dagstuhl möchte Klaus Tschira an dieser Stelle als Förderer und Gönner der Wissenschaft und vor allem als großzügigen Unterstützer von Schloss Dagstuhl und dblp würdigen.

Spender und Förderer

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

Neben zahlreichen Buchspenden durch Autoren und Gäste, erhielt das Zentrum einige Bildbände *Masters of*

sored “Year of Science 2014: Digital Society” initiative, and in the annual conference of the Leibniz Association in Berlin from November 26–28, 2014. Information about the center’s scientific mission and activities was distributed at both events.

Further details about public relations and professional training at Schloss Dagstuhl can be found in Chapter 5.

Joint Outreach with the Heidelberg Laureate Forum

In 2014, Schloss Dagstuhl was pleased again to cooperate with the Heidelberg Laureate Forum[2] (HLF), which involved the Dagstuhl Seminars “Neural-Symbolic Learning and Reasoning” (14381) and “Control Theory meets Software Engineering” (14382). We offered selected HLF 2014 participants – young researchers in the early phases of their careers – the opportunity to join the seminars as full participants during the week before the second annual edition of the forum on September 21–26, 2014. The HLF brings winners of the Abel Prize, the Fields Medal and the ACM Turing Award prize together with exceptionally talented young scientists from around the world. Satisfied with the outstanding success of the initiative, both partners agreed to renew the cooperation in 2015.

In preparation for this jointly sponsored week, Schloss Dagstuhl opened its doors on May 13, 2014 to a special visit from Dr. Klaus Tschira and several accompanying staff members from the Heidelberg team. Punctuated by Dr. Tschira’s thoughtful and friendly advice throughout the day, the collaborative exchange included a tour of the conference center, lunch and programmed round-table discussions with Scientific Director Prof. Raimund Seidel, former Scientific Director Prof. Reinhard Wilhelm, and core staff of Schloss Dagstuhl. The meeting proved to be Dr. Tschira’s last official visit to Schloss Dagstuhl before his unexpected death on March 31, 2015, which occurred during the writing of this report. Schloss Dagstuhl looks back on the day as a tribute to Klaus Tschira, a great sponsor and patron of science and, in particular, a generous supporter of Schloss Dagstuhl and dblp.

Sponsors and Donors

Schloss Dagstuhl is grateful to its scientific guests and institutional colleagues for generously donating funds to support its core scientific work and books to its research library.

In 2014, the center was glad to receive numerous private

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

Abstraction von der Klaus Tschira Stiftung. Dieser Bildband enthält Portraits von den noch lebenden Preisträgern des Turingpreises, des Abelpreises, der Fields-Medaille und des Nevanlinna-Preises. Die Portraits wurden von dem Fotografen Peter Badge erstellt, der von der Klaus Tschira Stiftung 2012 beauftragt wurde.

In 2014 erhielt die Dagstuhl-Bibliothek erneut zahlreiche Spenden für die Bibliothek in Form von Monographien von den Verlagshäusern Springer, O'Reilly, und Pearson, der European Association of Computer Graphics, und der Society for Industrial and Applied Mathematics. Insgesamt erhielt das Zentrum im Berichtszeitraum 949 Bände als Spenden, entsprechend 840 Monographien von den Verlagshäusern Springer im Wert von 58 527 €.

Wie in den Jahren 2011, 2012 und 2013 förderte die Klaus Tschira Stiftung 2014 erneut die Bibliographiedatenbank dblp mit einer Spende von 67 800 €. Dagstuhl Publishing – vor allem das Open-Access Journal *Leibniz Transactions on Embedded Systems* (LITES)[3] – erhielt 2014 substantielle Unterstützung durch Google Deutschland in Höhe von 10 000 US-Dollar (ca. 7 250 €).

Neben diesen Spenden durfte Schloss Dagstuhl im Jahr 2014 sich über eine dreidimensionalen Reproduktion von M. C. Eschers sogenannter „Belvedere" Zeichnung freuen, eine Schenkung von Prof. Gershon Elber von Technion für unsere permanente Kunstsammlung freuen. Prof. Elbers Forschungsgebiete umfassen unter anderem computergestützte, geometrische Konstruktion und Computergrafik. Mit Hilfe eines mehrlagigen Herstellungsverfahrens schuf Prof. Elber die außergewöhnliche Skulptur. Das Objekt steht momentan in einer Glasvitrine im Foyer des historischen Gebäudes von Schloss Dagstuhl.

Gründung des „Vereins zur Förderung von Schloss Dagstuhl"

Unterstützung durch private Gönner gab es seit Anbeginn des Zentrums. In 2014 wurde die bestehende Infrastruktur dafür auf eine neue Basis gestellt. Die Stiftung „Informatikzentrum Schloss Dagstuhl" zur Förderung von Wissenschaft und Forschung in Schloss Dagstuhl, an welche bisher Spenden gerichtet werden konnten, wurde neu aufgestellt durch Gründung eines Vereins. Förderer von Schloss Dagstuhl können nun Mitglied des gemeinnützigen „Vereins zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e. V." werden. Der Verein wurde im Juli 2014 gegründet und verwaltet seit dem die Mittel der bisherigen Stiftung, um Wissenschaft und Forschung in Schloss Dagstuhl zu fördern. Mehr Informationen sowie Ansprechpartner des Vereins sind über die Webseite des Vereins[4] erhältlich.

book donations from guests and partners, including several copies of the famous illustrated book *Masters of Abstraction* as a special gift from the Klaus Tschira Foundation. The book contains black and white photographic portraits of all living recipients of the Turing Award, the Abel Prize, the Fields Medal and the Nevanlinna Prize. The portraits were created by the photographer Peter Badge, who was sent by the Klaus Tschira Foundation on a worldwide journey to that end in 2012.

The center's scientific library also received in 2014 a large number of monographs from the Springer, O'Reilly, and Pearson publishing houses, the European Association of Computer Graphics, and the Society for Industrial and Applied Mathematics. The donated volumes totaled 949, including 840 monographs worth 58,527 € from the Springer publishing house.

With respect to grants, Schloss Dagstuhl was grateful to receive again in 2014 a grant of 67,800 € from Klaus Tschira Foundation to support the dblp computer science bibliography. The bibliography had already benefited from a generous donation for project work in 2011–2013. Dagstuhl Publishing – specifically the open-access journal *Leibniz Transactions on Embedded Systems* (LITES)[3] – also received strong external support in 2014 in the form of 10,000 US-Dollar (approx. 7,250 €) from Google Germany

In parallel to these donations, Schloss Dagstuhl received in 2014 as a gift for its permanent art collection a three-dimensional reproduction of M. C. Escher's so-called 'impossible' "Belveder" drawing from Prof. Gershon Elber of Technion. A computer scientist whose research interests include computer aided geometric designs and computer graphics, Prof. Elber created the unusual and beautiful sculpture thanks to a layered manufacturing process. The object currently stands in a glass vitrine in the foyer of the historic main building at Schloss Dagstuhl. Detailed information about the Schloss Dagstuhl art donation program can be found in Chapter 10.

Founding of the Association "Friends of Schloss Dagstuhl"

Private support for Schloss Dagstuhl has always been welcome. In 2014, the existing infrastructure for this was put on a new footing as the Dagstuhl Foundation, which formerly had received direct donations in support of the center's core work, was reorganized and a new association was formed. Persons who wish to further the basic mission and goals of the center now have the option to join the nonprofit association "Friends of Schloss Dagstuhl" (*Verein zur Förderung von Schloss Dagstuhl – Leibniz-Zentrum für Informatik e. V.*). Founded in July 2014, the association administers the goods and resources of the Dagstuhl Foundation (*Stiftung Informatikzentrum Schloss Dagstuhl*) to support science and research at Schloss Dagstuhl. Detailed information and contact persons can be found under the Friends of Schloss Dagstuhl website[4].

[3] http://www.dagstuhl.de/lites/
[4] http://www.friends-of-dagstuhl.de

Auszeichnungen

Im August 2014 erreichte Schloss Dagstuhl die Nachricht, dass der 2012 in Betrieb genommenen Neubau des Gästehauses mit dem Bauherrenpreis der Architektenkammer des Saarlandes[5] ausgezeichnet werden wird. Das Gästehaus wurde von dem Berliner Architekturbüro av-a Veauthier Architekten entworfen. Der Preis wurde nur an 6 weitere der insgesamt 57 eingereichten Projekte verliehen. In seiner Einreichung erläuterte das Architekturbüro, Herausforderung und Ziel seien gewesen, den einzigartigen Charakter und die Geschichte der Schlossanlage durch den Neubau nicht umzuformulieren. Stattdessen werde mit dem Garten als großem, natürlichem und umgrenztem Zwischenraum ein visueller und konzeptueller Kontrapunkt zum Schlossgebäude geschaffen. Dr. Lindig, damaliger Technisch-administrativer Direktor von Schloss Dagstuhl, war als Bauherr während der Entwurfs- und Bauphase bis zur Einweihung des Gebäudes 2012 mit der Abstimmung zwischen Schloss Dagstuhl und dem Architekturbüro betraut. Er nahm am 1. Oktober 2014, zusammen mit Frau Meißner und Prof. Seidel, den Preis entgegen. Der Preis war Anlass für eine Sonderausstellung in den Räumen des Saarländischen Rundfunks, einen Film und ein Buch, „Bauherrenpreis 2014 der Architektenkammer des Saarlandes“, in denen das Gästehaus zusammen mit den anderen ausgezeichneten Arbeiten vorgestellt wird.

Awards

Schloss Dagstuhl was pleased to learn in August of 2014 that it had been awarded the Bauherrenpreis of the Saarland Chamber of Architects[5] for its modern guest house, designed by the Berlin-based architectural firm Veauthier Meyer Architects. The prize was one of only seven awarded out of competing 57 entries. Describing the challenges of the project, the architectural firm had noted in its award-winning contest submission that the goal was not to break with the unique character and history of the grounds at Schloss Dagstuhl, but rather to create a visual and conceptual counterpoint to the monumental “Schloss” from a point of relative seclusion across the long, natural space of the Schloss garden. Schloss Dagstuhl Scientific Director Prof. Seidel, Technical Administrative Director Heike Meißner, and former Technical Administrative Director Dr. Lindig, who had facilitated the arrangements with the architects during the building’s planning and construction phases prior to its inauguration in 2012, received the award at a ceremony in Saarbrücken on October 1, 2014. The award gave rise to a special exhibit, film and book, “Bauherrenpreis 2014 der Architektenkammer des Saarlandes,” in which the guest house was featured along with the other awarded designs.

Gästeservice und Einrichtungen

Schloss Dagstuhl bietet allen Teilnehmern von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops, die mit Kindern anreisen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser Service kann im voraus gebucht werden. Alternativ ist es Eltern auch möglich eine Begleitperson zur Betreuung des Kindes oder der Kinder mitzubringen. Die Kosten für Verpflegung und Unterkunft für die Begleitperson als auch der Kinder übernimmt Schloss Dagstuhl. Im Jahre 2014 wurden 22 Kinder auf Schloss Dagstuhl betreut. Davon wurden 9 Kinder durch einen Tagesmutter und 13 weitere durch Verwandte betreut. Positive Rückmeldungen von den Eltern sowie Dankeschön-Briefe und Bastelarbeiten, die die kleinen Gäste 2014 für das Kindergästebuch von Schloss Dagstuhl angefertig haben, sind in Kapitel 3 und über den ganzen Bericht verteilt zu finden.

2014 wurden Wartungarbeiten durchgeführt, die garantieren dass die Sicherheitseinrichtungen auf Schloss Dagstuhl auch weiterhin dem neuesten Stand entsprechen, um Gästen auf Schloss Dagstuhl die größtmögliche Sicherheit bieten zu können. Im Anschluss an die 2013 ausgeführten Vorarbeiten bezüglich des Feuer- und Warnsystems wurden 2014 die Aufzüge in allen drei Gebäuden mit einem 24-Stunden-Notrufdienst ausgestattet. Daneben wurden neue Fluchtwegpläne und Beschilderungen für den Brandfall in deutscher und englischer Sprache ausgehängt,

Guest Services and Facilities

Schloss Dagstuhl gladly offers to organize child care with a certified nanny for participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops who need to visit our center with young children. The service, which supports families and particularly women computer scientists, can be booked in advance of the seminar. Parents also have the option to bring along their own “nanny,” usually a spouse or relative, whose room and board costs are absorbed by the center just as they are for children. In 2014, Dagstuhl hosted 22 children, 9 of whom were cared for by a nanny on site and 13 by relatives. Positive feedback from parents, and thank-you notes and artwork created by center’s youngest guests in 2014 for the Dagstuhl children’s guest book, can be found in Chapter 3 and throughout this report.

Facility updates in 2014 were relatively minor and focused on ensuring the continued maximum safety and security of all guests at Schloss Dasgstuhl. Following up on preparations carried out in 2013 with respect to the center’s main fire alert and response system, the elevators in all three buildings were outfitted in 2014 with a 24-hour emergency alert service. New fire escape maps and place signs were also installed to reflect recent changes in the grounds and offer bilingual information in English and German. Thanks in part to its unusual construction features, Schloss Dagstuhl is particularly well suited for fire drills and offers its grounds for this purpose once a year. The drills, carried

[5] http://www.aksaarland.de/bauherren/bauherrenpreis-2014-die-preistraeger

die die jüngsten Veränderungen der Schlossanlage berücksichtigen. Neben diesen technischen Errungenschaften profitiert Schloss Dagstuhl auch von jährlichen Feuerübungen, die die Sicherheit des Warnsystems verstärken.

Im Jahre 2013 begann die Modernisierung der sanitären Anlagen sowie der Ausstattungen der Bäder in den Gästezimmern des Neubaus. Das Projekt wurde im folgendem Jahr mit der vollständigen Renovierung von fünf Badezimmern fortgesetzt. Die Renovierungsarbeiten in dem seit 24 Jahren bestehenden Gebäude werden auch 2015 fortgesetzt werden. Im Rahmen dessen ist in Planung, die Fenster zum Innenhof zu erneuern. Die öffentlich zugänglichen Bereiche profitierten ebenfalls von einigen kleineren Veränderungen, die die Funktionalität, die Zugänglichkeit und den Entwurf der Außenanlage des Konferenzzentrums erhalten, wie zum Beispiel die Einrichtung eines Behindertenparkplatzes oder ein neuer Parkettfussboden aus Eiche in einem Konferenzraum des Schlosses. Ein kleines Detail von besonderem Interesse für unsere Gäste ist die sogenannte „White Wall“ – eine Wand des Raums die vollständig als White-Board genutzt wird – im Konferenzraum Trier, der 2014 mit einer digitalen Kamera und besonderer Software ausgestattet wurde. Seminarteilnehmer können die Wand nach einer Diskussion abfotografieren und die Ergebnisse an eine gewünschte E-Mail-Adresse schicken. Entwickelt hat dieses Format Steven Schrot, der in der IT-Abteilung tätig ist.

Die Wünsche und Vorstellungen unserer Gäste haben einen großen Einfluss auf die Gestaltung und Benutzerfreundlichkeit unseres Programms. Ein Beispiel ist der „Dagstuhl Timer“, der bei der Koordinierung von Vorträgen behilflich sein kann; die Anwendung wurde von Prof. Matt Duckham entwickelt und ist über die Schloss Dagstuhl Webseite[6] für die Seminar-Organisatoren als Hilfsmittel einsetzbar.

out by the local fire brigade, contribute to the overall safety of the center's emergency response system.

Since 2013, the center has also undertaken to fully modernize the installations and retrofit guest room bathrooms in the New Building, a project that went forward in 2014 with the complete renovation of five bathrooms. These ongoing improvements in the New Building, which is now over 24 years old, are expected to continue in 2015 with the replacement of the building windows along the inner courtyard. Common spaces also benefitted from a few minor interventions to maintain the functionality, accessibility and design of the conference center grounds, such as the creation of a disabled parking space and the installation of new oak flooring in one of the Schloss meeting rooms. One small detail of special interest to seminar guests is the fact that the so-called “White Wall” in the Trier conference room was also outfitted in 2014 with a built-in digital camera with special software. Seminar participants can use the camera to photograph the wall following a discussion, and send the resultant images directly to their email inboxes. The software and construction was developed by Schloss Dagstuhl IT staff member Steven Schrot.

Ideas from guests often bring about lasting improvements in the quality of the program experience at Schloss Dagstuhl, and 2014 offered fresh opportunities to witness this creativity in action. One example is the “Dagstuhl Timer,” a software tool designed to help coordinate talks during scientific meetings; the timer was created by Prof. Matt Duckham and is now referenced on the Schloss Dagstuhl webpages[6] as a seminar planning tool.

[6] http://www.dagstuhl.de/en/program/dagstuhl-seminars/info-for-organizers/

2 Das wissenschaftliche Programm 2014
The 2014 Scientific Program

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 in der Informatik-Community 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 wenn nicht sogar erschöpfende Übersicht über die Gebiete der Informatik, die derzeit weltweit diskutiert werden.

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 zwölf Ausgaben veröffentlicht. Jede Ausgabe dokumentiert jeweils die Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops eines Monats. Die Dagstuhl Reports sind über die Dagstuhl-Website frei zugänglich.[7]

In den Kapiteln 4 und 13 sind alle Veranstaltungen, die 2014 auf Schloss Dagstuhl stattfanden, aufgelistet, zusammen mit Zusammenfassungen der Seminare und Perspektiven-Workshops. Auf der Dagstuhl-Website ist das 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 if not comprehensive view of the areas currently under discussion in the international 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 twelve issues per year; each issue documents the Dagstuhl Seminars and Dagstuhl Perspectives Workshops of a given month. Dagstuhl Reports are open-access and can be downloaded at from the Dagstuhl website.[7]

Chapters 4 and 13 provide a comprehensive list of all events that took place at Schloss Dagstuhl during the year under review and summaries of the 2014 Seminars and Perspectives Workshops. A program covering the coming 24 months is available on the Dagstuhl website.

[7] 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 2014 statt fanden, sind in Fig. 2.1 aufgelistet.

Die Ergebnisse der intensiven Diskussionen werden in einem Manifest zusammengefasst, welches die offenen Probleme und die möglichen Forschungsperspektiven für die nächsten 5–10 Jahre aufzeigt. Dagstuhl koordiniert die gezielte Weitergabe dieses Manifests, um forschungsspezifische Impulse an deutsche und europäische Institutionen der Forschungsförderung zu geben (EU, BMBF, DFG, etc.). Kurzfassungen der Manifeste werden regelmäßig im Forum des *Informatik Spektrum* (Springer-Verlag) vorgestellt. Die vollständigen Manifeste werden in unserer Fachzeitschrift Dagstuhl Manifestos[8] veröffentlicht.

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

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 2014 are listed in Fig. 2.1.

The results of the in-depth discussions of each workshop are presented in a manifesto detailing open issues and possible research perspectives in that specific field for the coming 5–10 years. Schloss Dagstuhl coordinates the targeted dissemination of this manifesto as research policy impulses to German and other European research donors and sponsors (EU, German Federal Ministry of Education and Research, DFG, etc.). Short versions of the manifestos are regularly presented in a forum of the *Informatik Spektrum* journal (published by Springer); full versions of the manifestos are published in our periodical Dagstuhl Manifestos[8].

Past and upcoming Dagstuhl Perspectives Workshop can be found on our web site.[9]

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

Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing http://www.dagstuhl.de/14022
Exploring Interdisciplinary Grand Challenges in ICT Design to Support Proactive Health and Wellbeing http://www.dagstuhl.de/14272
Categorical Methods at the Crossroads http://www.dagstuhl.de/14182
Massive Open Online Courses: Current State and Perspectives http://www.dagstuhl.de/14112
Privacy and Security in an Age of Surveillance http://www.dagstuhl.de/14401

Fig. 2.1
Dagstuhl Perspectives Workshops held in 2014.

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 2014 wieder ein hohes Interesse am Organisieren von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops durch die Einreichung von insgesamt 98 Anträgen in den Antragsrunden im Januar und Juni 2014. Etwas mehr als 66 % der eingereichten Anträge wurden genehmigt. Der Großteil der Anträge genügen den Antragskriterien überdurchschnittlich. Ablehnungen gibt es zum Beispiel vor dem Hintergrund, dass ein thematisch ausgeglichenes Seminarprogramm aus den eingegangenen Anträgen erstellt werden muss und daher in Themenbereichen mit einer Häufung von Anträgen nur die – im Vergleich – interessantesten Anträge angenommen werden können. In den vergangenen 7 Jahren variierte die Rate der angenommen Anträge zwischen 66 % und 77 % (siehe Fig. 2.3).

Unter den 65 in 2014 neu genehmigten Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops gab es auch dieses Mal wieder verschiedene Konstellationen bzgl. Dauer und Größe (vgl. Fig. 2.2). Einige dieser Seminare konnten sogar bereits 2014 ausgerichtet werden, der Großteil wurde jedoch für das Programm in 2015 und zu Beginn von 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 2014, submitting 98 proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops during the January 2014 and June 2014 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 77 % (see Fig. 2.3).

Among the 65 Dagstuhl Seminars and Dagstuhl Perspectives Workshops accepted in 2014 there is – as usual – a wide variation with regard to length and size (see Fig. 2.2). Most of these seminars are part of the 2015 and early 2016 program, although it was possible to schedule some of them already in 2014.

	small	large
short	12	1
long	19	33

Fig. 2.2
Number of Dagstuhl Seminar and Dagstuhl Perspectives Workshop proposals accepted by the Scientific Directorate in 2014.
Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

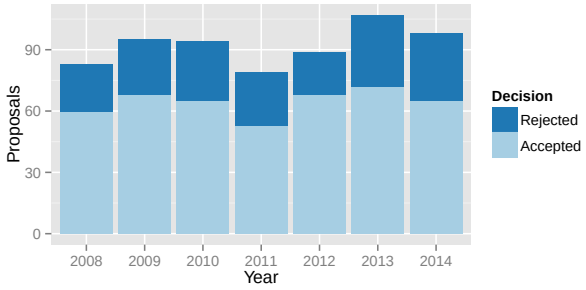

Fig. 2.3
Overview of proposed and accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2008–2014.

Dagstuhl-Seminarprogramm 2014

2.4 *The Dagstuhl Seminar Program in 2014*

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 zusammengelegt. In 2014 fanden in 33 von 48 Wochen Seminare parallel statt, was 69 % der verfügbaren Wochen entspricht. Durch einen Vergleich mit den Zahlen aus 2013 (62 %), 2012 (35 %) und 2011 (12 %) kann man feststellen, dass derweil eine gewisse Sättigung eingetreten ist und man sich nahe an der optimalen Auslastung befindet. In Fig. 2.4 ist die Entwicklung der vergangenen Jahre dargestellt.

Durch die Möglichkeit, zwei Seminare parallel in einer Woche zu veranstalten, ist in den letzten Jahren auch die Gesamtanzahl an Seminaren pro Jahr gestiegen. Mit 75 Seminaren in 2014 bewegt man sich mittlerweile am maximal Möglichen.

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. In 2014, there were two parallel seminars per week on 33 out of 48 weeks, or on 69 % of all available weeks. This rate indicates a near-optimal saturation state, as compared against the figures 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 75 in 2014 – its near-maximum.

Angaben zu Teilnehmern und Organisatoren

2.5 *Participant and Organizer Data*

Die Teilnehmer der Seminare kommen aus aller Welt und viele 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 2014 circa die Hälfte (49 %, 1 257 von 2 590) der Gäste das erste Mal an einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop teil, während 21 % der Teilnehmer an nur einem vorherigen Seminar teilgenommen hatten (siehe Fig. 2.5a).

Ein beträchtlicher Anteil der Gäste besteht aus jungen Wissenschaftlern, die am Anfang ihrer Karriere stehen, und für die der Aufenthalt in Dagstuhl oftmals prägend

Participants in Dagstuhl Seminars come from all over the world and a significant number of them choose to repeat the experience. Nevertheless, we see many fresh new faces every year, reflecting the changing informatics research across the globe. As in the previous year, also in 2014 about the half (1,257 of 2,590, or nearly 49 %) of Dagstuhl Seminar participants were first-time visitors to Dagstuhl, followed by 538 participants (nearly 21 %) who had already attended one previous seminar at the center (see Figure 2.5a).

A healthy number of these guests were young researchers at the start of their careers, for whom the

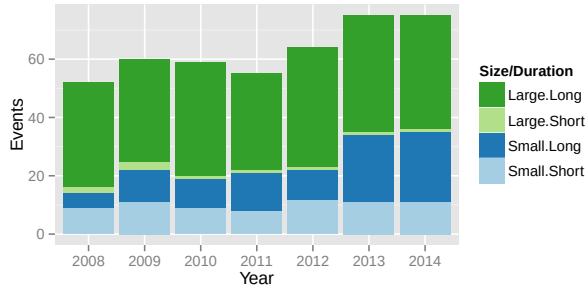

Fig. 2.4
Size and duration of Dagstuhl Seminars and Dagstuhl Perspectives Workshops held in 2008–2014. Small = 30-person seminar, large = 45-person seminar, short = 3-day seminar, long = 5-day seminar.

ist für den weiteren Verlauf ihres Lebenswegs. Etwa 27 % der Gäste der Seminare und Workshops in 2014, die an unserer Umfrage zur Qualitätskontrolle teilgenommen haben, stuften sich selbst als Nachwuchswissenschaftler ein (siehe Fig. 2.5b). Diese ausgewogene Verteilung zwischen Nachwuchswissenschaftlern und erfahrenen Forschern ist im Laufe der Jahre relativ konstant geblieben, was die Bemühungen des Zentrums zur Aufrechterhaltung der „Dagstuhl-Verbindung“ zwischen herausragenden jungen Wissenschaftlern und ihren erfahrenen Kollegen zeigt.

Mit 76 % war der Anteil von Gästen aus dem Ausland 2014 erneut sehr hoch. Das Diagramm in Fig. 2.5c zeigt die regionale Verteilung der Gäste für 2014 bei Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops. Mehr Details kann Kapitel 12 entnommen werden.

In 2014 war fast die Hälfte aller Organisatorenteams des Seminar-Programms hinsichtlich des Geschlechts gemischt, wobei dieses Verhältnis im Vergleich zu den Vorjahren konstant geblieben ist (siehe Fig. 2.6a). Der Anteil an weiblichen Seminarteilnehmern war mit 15,7 % wieder erfreulich hoch (siehe Fig. 2.6b).

Dagstuhl experience can be of lifelong value. Approximately 27 % of 2014 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 76 %, the proportion of seminar and workshop guests with a non-German affiliation in Dagstuhl Seminars was extremely high again during 2014. The chart in Fig. 2.5c shows the regional distribution of our Dagstuhl Seminar and Dagstuhl Perspectives Workshop guests in 2014. For a detailed breakdown please refer to Chapter 12.

In 2014, nearly half of all organizer teams in our scientific program were mixed with respect to gender, a proportion that has remained relatively unchanged in comparison to most previous years since 2008 (see Fig. 2.6a). The percentage of female seminar participants was also high both in total and relative terms, at 15.7 % (see Fig. 2.6b).

2.6 Themen und Forschungsgebiete / *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 Programm zu geben. Kapitel 4 bietet mit den Kurzzusammenfassungen der Seminar- und Workshops einen vollständigen Überblick über das wissenschaftliche Programm des Jahres 2014.

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 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 2014. 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 program of 2014. The seminar summaries in Chapter 4 provide a full overview of the program.

Some of the 2014 seminars looked at issues that have an impact on society and dealt critically with informatics as a scientific discipline and its influence (14042 – *Do-it-yourself Networking: an Interdisciplinary Approach*; 14052 – *Ethics in Data Sharing*; 14112 – *Massive Open*

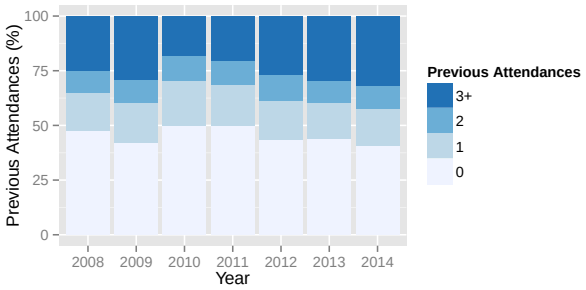

(a)Distribution of the number of previous attendances of participants, according to survey data.

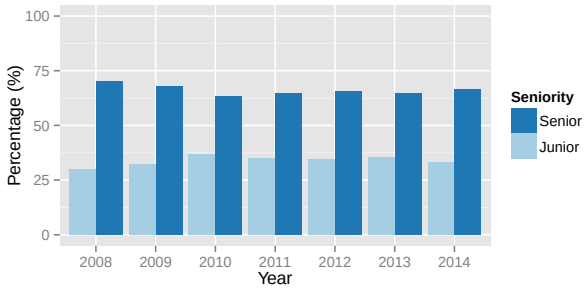

(b)Percentage of junior researchers, according to survey data.

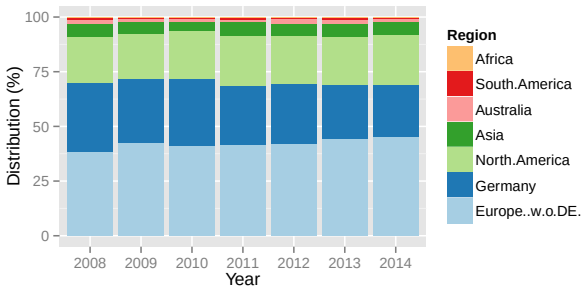

(c)Distribution of the origin of participants by region.

Fig. 2.5
Participants of Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2008–2014.

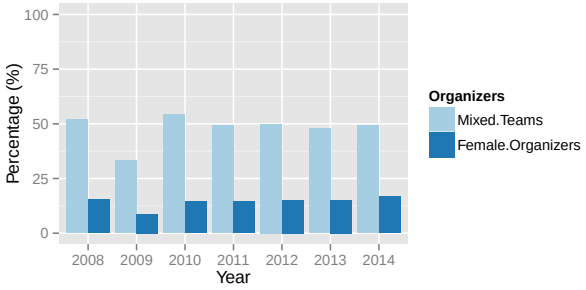

(a)Female organizers and mixed-gender organizer teams.

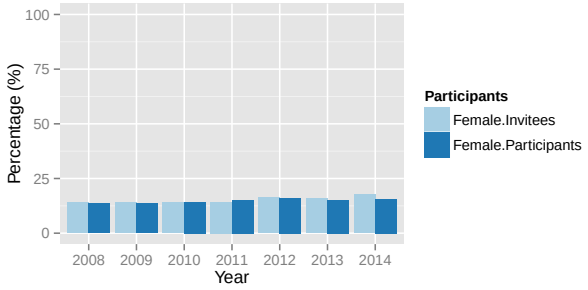

(b)Female invitees and participants.

Fig. 2.6
Female researchers at Dagstuhl Seminars and Dagstuhl Perspectives Workshops in 2008–2014.

Einige der Seminare in 2014 hatten gesellschaftlich relevante Themen zum Inhalt, die die Informatik als wissenschaftliche Diszplin und ihre Rolle bezüglich den Auswirkungen kritisch hinterfragt haben (14042 – *Do-it-yourself Networking: an Interdisciplinary Approach*; 14052 – *Ethics in Data Sharing*; 14112 – *Massive Open Online Courses: Current State and Perspectives*; 14401 – *Privacy and Security in an Age of Surveillance*; 14471 – *Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment*). Desweiteren gab es spannende Seminare, die eine wichtige Brücke zwischen den Geisteswissenschaften und der Informatik geschlagen haben (14301 – *Computational Humanities – Bridging the Gap Between Computer Science and Digital Humanities*; 14302 – *Digital Palaeography: New Machines and Old Texts*). Interessant erscheint auch, dass verstärkt kollektive Phänomene und Ressourcen wissenschaftlich

Online Courses: Current State and Perspectives; 14401 – *Privacy and Security in an Age of Surveillance*; 14471 – *Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment*). There were also illuminating seminars that built an important bridge between the humanities and informatics (14301 – *Computational Humanities – Bridging the Gap Between Computer Science and Digital Humanities*; 14302 – *Digital Palaeography: New Machines and Old Texts*). Interestingly, collective phenomena and resources are increasingly in the focus of scientific investigations (14512 – *Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis*; 14282 – *Crowdsourcing and the Semantic Web*). On the political level, critical infrastructures are of high interest currently, and they were of major priority for scientific analyses also in some seminars (14031 – *Randomized Timed and Hybrid Models for Critical Infrastructures*,

untersucht wurden (14512 – *Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis*; 14282 – *Crowdsourcing and the Semantic Web*). Kritische Infrastrukturen, die aktuell auf politischer Ebene ganz hoch im Kurs stehen, waren ebenfalls Ziel von Diskussionen und wissenschaftlicher Untersuchungen im Seminarprogramm (14031 – *Randomized Timed and Hybrid Models for Critical Infrastructures*, 14292 – *Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures*).

Generell umfasste das Seminar-Programm 2014 wie üblich eine breite Palette von Forschungsgebieten, die oftmals interdisziplinär an den Seminaren und Workshops beteiligt waren.

14292 – *Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures*).

In general, Dagstuhl's 2014 seminar program included the usual broad range of research areas, which often were part of interdisciplinary seminars and workshops.

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
- Sommerschulen, Weiterbildungsveranstaltungen, Lehrerfortbildungen, Ausbildung von jungen Journalisten und Volontären
- Klausurtagungen von Graduiertenkollegs, GI-Fachgruppen und anderen akademischen und industriellen Arbeitsgruppen
- in geringem Umfang internationale Informatik-Fachtagungen
- 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.

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
- 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
- a small number of international informatics conferences
- 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.

2.8 Qualitätssicherung

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

2.8 *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

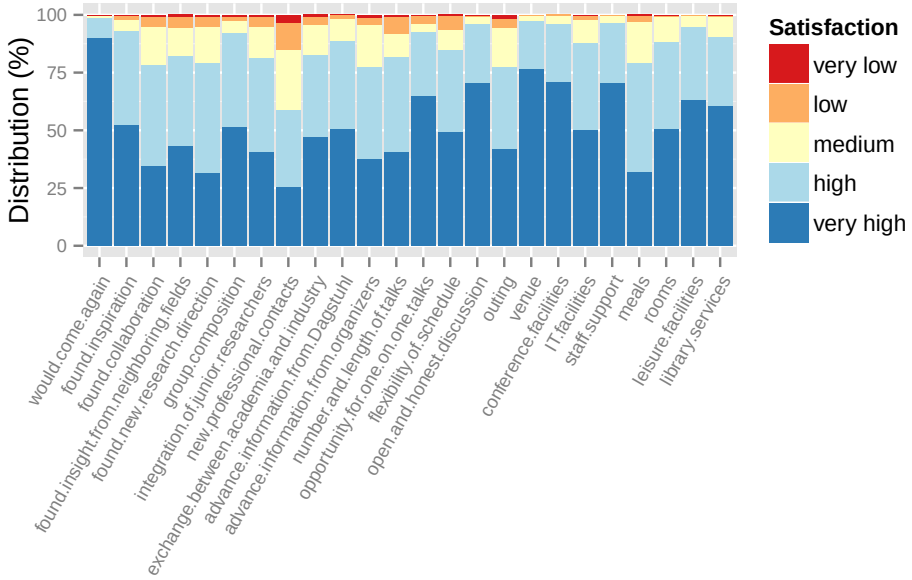

Fig. 2.7
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants in 2014. According to survey results.

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 2014 zu ausgewählten Aspekten ihres Aufenthaltes. Grundlage ist die Auswertung von 1 457 Fragebögen, welche die Meinung von etwa 56 % der 2 590 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.

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 1,457 questionnaires, representing the responses of about 56 % 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 *Utilization of the Center*

Auch 2014 konnte Schloss Dagstuhl die durch das neue Gästehauses ermöglichte hohe Auslastung weitgehend halten. Insgesamt gab es 2014 mit 13 522 Übernachtungen, wobei 11 414 Übernachtungen auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops entfielen. Bezogen auf die Seminar- und Workshopgäste bedeutet dies ein ähnliches Niveau wie in 2013. Es fanden im Berichtsjahr 112 Veranstaltungen mit insgesamt 3 409 Gästen statt. Weitere Details können Kapitel 12 entnommen werden.

Die Wochenenden blieben 2014 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 2014 einschließlich Dagstuhl-Se-

Thanks to the new guest house, Schloss Dagstuhl was able to uphold the high capacity utilization again in 2014. There were 13,522 overnight stays in total, with 11,414 overnight stays in seminars and perspective workshops. The latter was on a similar level as in 2013. The center hosted a total of 112 events with 3,409 guests in 2014. See Chapter 12 for further details.

Weekends were kept free in 2014, 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 2014, including Dagstuhl Seminars, Dagstuhl Perspectives Workshops, GI-Dagstuhl Seminars, and host-only events such as meetings and summer schools can

minaren, Dagstuhl-Perspektiven-Workshops, GI-Dagstuhl-Seminaren und Veranstaltungen (z.B. Sommerschulen), bei denen Schloss Dagstuhl nur Veranstaltungsort war, findet sich in Kapitel 13. Auf unserer Webseite ist ein Kalender[10] verfügbar, in welchem die anstehenden Veranstaltungen eingesehen werden können, ebenso wie weitere Informationen und Materialien zu allen vergangenen, aktuellen und zukünftigen Veranstaltungen.

be found in Chapter 13. See the Schloss Dagstuhl website to view our calendar[10] 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.

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

3 Resonanz
Feedback

Resonanz von Seminarteilnehmern

3.1

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 Daumer (Trium Analysis Online GmbH – München, DE)
14062 – The Pacemaker Challenge: Developing Certifiable Medical Devices | Dagstuhl Seminar | http://www.dagstuhl.de/14062

It was a great pleasure and honor for me to have been invited to contribute. [...] The location and infrastructure is outstanding; the spectrum of experience, know-how and interests of the participants was just right for the focus topic of the meeting; the "agile" way of planning and conducting the seminar is excellent.

Alan Wassyng (McMaster University – Hamilton, CA)
14062 – The Pacemaker Challenge: Developing Certifiable Medical Devices | Dagstuhl Seminar | http://www.dagstuhl.de/14062

Thank you for the support – Dagstuhl is a great experience, and a truly unique and wonderful institution!

Barbara Endicott-Popovsky (University of Washington – Seattle, US)
14092 – Digital Evidence and Forensic Readiness | Dagstuhl Seminar | http://www.dagstuhl.de/14092

The experience at Dagstuhl (my first) surpassed my expectations. It was a powerful opportunity to collaborate with colleagues and extend my research. I can't thank the Institute enough for providing this venue!

Hein Venter (University of Pretoria – Pretoria, ZA)
14092 – Digital Evidence and Forensic Readiness | Dagstuhl Seminar | http://www.dagstuhl.de/14092

Dagstuhl has been an incredible experience to me. The minds that have crossed my path have left huge footprints in my career, but more importantly, in my heart. It was, and is, a fantastic pleasure to work with you all. Your input has filled my career's gas tank with wonderful zeal.

Omri Weinstein (Princeton University – Princeton, US)
14121 – Computational Complexity of Discrete Problems | Dagstuhl Seminar | http://www.dagstuhl.de/14121

I chose this option for reporting my experience at Dagstuhl (complexity seminar). As a graduate student, it was my first time in Dagstuhl and definitely one of the best academic experiences I've ever had. Both the academic program and the organization in the facility were fantastic, and I felt interactions with the members of the seminar was highly productive. I hope to have the opportunity to be invited again. A special thanks and appreciation to the organizing team and the staff at Dagstuhl!

Toyoaki Nishida (Kyoto University – Kyoto, JP)
14131 – Computational Models of Cultural Behavior for Human-Agent Interaction | Dagstuhl Seminar | http://www.dagstuhl.de/14131

My experience at Schloss Dagstuhl was really intriguing; it was very productive and fruitful as I came up with many new thoughts from the discussions with other participants. Supporting people were nice and kind, and the place was excellent for intimate discussions. In particular, I liked the academic and friendly atmosphere very much. Thank you so much for having supported our event. Definitely I look forward to coming back in the near future.

Kathi Fisler (WPI – Worcester, US) and Shriram Krishnamurthi (Brown University – Providence, US)

We are both computer science professors who have been fortunate to visit Dagstuhl several times. Because we collaborate on some of our research, we have sometimes been invited to the same event. After we became parents we worried about how we might handle such a joint invitation, and were delighted to find out about Dagstuhl's child-care option.

Having now experienced it multiple times, we can safely say the Dagstuhl option should be regarded as the gold standard for parental and child support. The Dagstuhl staff are unfailingly helpful, the sitters have been excellent people, and the location is perfect to accommodate this arrangement. Relative to the child-care programs we have experienced at some conferences [...] the Dagstuhl option is beyond compare. Finally, the cost is extremely reasonable, especially for the level of service. Without the child-care option, it is unclear whether the two of us would have been able to attend week-long events together. With it, we are able to come stress-free.

Ultimately, the Dagstuhl program saves parents from having to make a choice as to who can exploit the unique opportunity that a Dagstuhl provides, and thus undoubtedly helps female participants, who may be more likely to otherwise be saddled with parenting duties. It also sends a message that top-caliber research and family are not incompatible. Traveling with children is both stressful and rewarding; Dagstuhl does its best to minimize stress and maximize reward. Overall, it therefore makes our discipline a more civilized one.

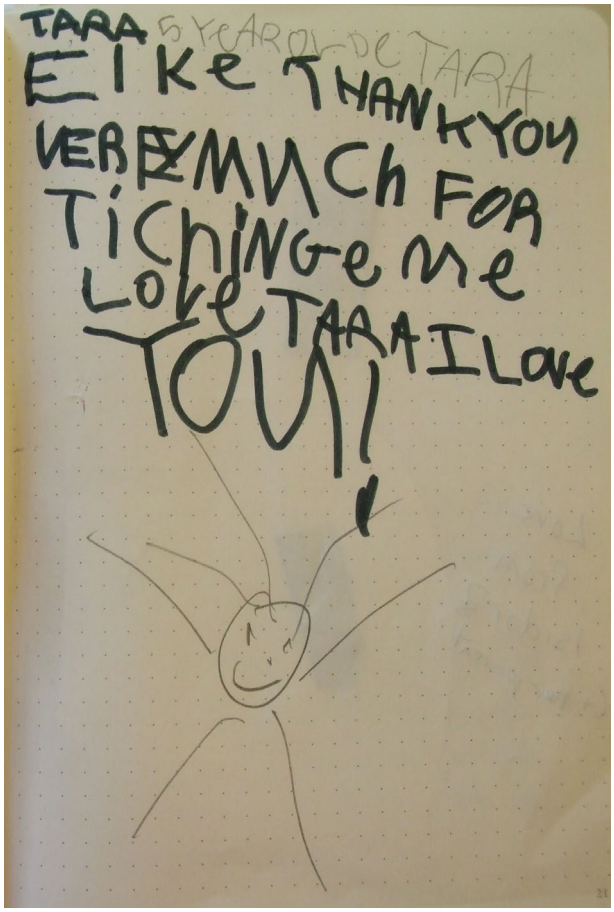

Fig. 3.1
Drawing for the Dagstuhl children's guest book by Tara (5), daughter of Dagstuhl Seminar 14281 participants Kathi Fisler and Shriram Krishnamurthi.

Charles W. Krueger (BigLever Software – Austin, US)

14172 – Unifying Product and Software Configuration | Dagstuhl Seminar | http://www.dagstuhl.de/14172

Thank you, I look forward to attending this Dagstuhl Seminar and revisiting one of the most wonderful research havens on the planet.

Michael Fellows and Frances Rosamond (Charles Darwin University – Darwin, AU)

14309 – Research Stay | Meeting | http://www.dagstuhl.de/14309

Thank you VERY VERY much! We are staying in Europe so long because we have been involved as organisers of various meetings in Europe this year [...] It is just too far to Australia to go back and forth in the gaps, and being at Dagstuhl is just perfect. We truly appreciate your hospitality, and the wonderful library here, etc.

Sani Abba (Chosun University – Gwangju, KR)

14382 – Control Theory meets Software Engineering | Dagstuhl Seminar | http://www.dagstuhl.de/14382

I write to thank you all for the great opportunity you gave me to participate in the 2nd Heidelberg Laureate Forum and the Dagstuhl GI seminar 14382 "Control theory meets software engineering." As a young PhD candidate and a researcher, I learnt a lot. Thank you once again for organizing this wonderful intellectual gathering and professional collaboration and networking.

Michael Witt (Purdue University – West Lafayette, US)

14432 – DDI Lifecycle: Moving Forward | Dagstuhl Seminar | http://www.dagstuhl.de/14432

It was a pleasure to spend time in your library and also to talk with you [Schloss Dagstuhl librarian Petra Meyer] last week during my seminar at Dagstuhl. I have collection development responsibilities for CS at Purdue, and your library gave me some ideas for improving our services and collection. You have done an outstanding job.

3.2 Resonanz unserer Seminarorganisatoren

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 folgende geben mit freundlcher Genehmigung der Autoren einige der Kommentare unsere Seminarorganisatoren wieder.

3.2 *Feedback from Seminar Organizers*

The success of Schloss Dagstuhl depends to a large extent on our outstanding seminar organizers, who continually enrich the scientific program with a range of interesting and new topics. We are very glad to be able to provide services and an environment that organizers appreciate. The following comments from organizers are excerpted from the Dagstuhl Report for each of the cited seminars.

Organizers of Dagstuhl Seminar 14062

14062 – The Pacemaker Challenge: Developing Certifiable Medical Devices | Dagstuhl Seminar | http://www.dagstuhl.de/14062

This new funding effort would not have been possible without the possibility to meet at Dagstuhl several times to find common ground and determine an overall strategy. In short, we were very happy with the discussions, work and impact of the Dagstuhl seminar on translation to morphologically rich languages. [...] Finally, we would like to once again thank the staff of Dagstuhl for facilitating these unique scientific discussions which we are confident will have a strong impact on future research on the important problem of statistical techniques for translation to morphologically rich languages.

Fig. 3.2
„Elizabeth Churchill who was the amazing organizer of the excellent Dagstuhl Workshop signing her books in Dagstuhl." Twitter post by 14272 Dagstuhl Seminar participant Adrian David Cheok.

Organizers of Dagstuhl Seminar 14082

14082 – Visualization and Processing of Higher Order Descriptors for Multi-Valued Data | Dagstuhl Seminar | http://www.dagstuhl.de/14082

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

Organizers of Dagstuhl Seminar 14111

14111 – Combinatorics and Algorithmics of Strings | Dagstuhl Seminar | http://www.dagstuhl.de/14111

Such a unique assembly of major experts in word equations and their contributions at Dagstuhl was rather unique and a remarkable event. In the light of such developments, it can be safely claimed that this seminar was a great success. Given the quality of presentations on this seminar and the constructive intensity of discussions, it is self-evident that a follow-up should be organized. We are grateful to all participants for their contributions to this successful seminar as well as to the staff of Schloss Dagstuhl for their great service.

Organizers of Dagstuhl Seminar 14121

14121 – Computational Complexity of Discrete Problems | Dagstuhl Seminar | http://www.dagstuhl.de/14121

We like to thank the staff at Dagstuhl who – as usual – provided a marvelous surrounding to make this a successful meeting with ample space for undisturbed interactions between the participants.

Organizers of Dagstuhl Seminar 14221

14221 – Geometric Modeling | Dagstuhl Seminar | http://www.dagstuhl.de/14221

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

Organizers of Dagstuhl Seminar 14201

14201 – Horn formulas, directed hypergraphs, lattices and closure systems: related formalisms and applications | Dagstuhl Seminar | http://www.dagstuhl.de/14201

We believe that the seminar was very successful in bringing together a critical mass of researchers from different communities and in providing a platform for personal contacts and scientific interchange between the participants. [...] Besides presentations, the program offered room for open discussions and informal working groups. As a major outcome, a special issue of the journal *Theoretical Computer Science*, co-edited by the organizers, will be devoted to the themes of the seminar. We hope this could serve as a reference material for future interdisciplinary research in the field. Schloss Dagstuhl and its staff provided a very convenient and stimulating environment. The seminar participants appreciated the cordial atmosphere which improved mutual understanding and inspiration. The organizers of this seminar wish to thank all those who helped to make the seminar a fruitful research experience.

Organizers of Dagstuhl Seminar 14382

14382 – Control Theory meets Software Engineering | Dagstuhl Seminar | http://www.dagstuhl.de/14382

We finally had the time to evaluate, also with our invitees, the seminar. It seems it was a success, and this is also because of the thorough, patient, and friendly support you gave us during the organization and our staying there. We really wanted to thank you for your professionalism and availability!

Organizers of Dagstuhl Seminar 14421

14421 – Optimal Algorithms and Proofs | Dagstuhl Seminar | http://www.dagstuhl.de/14421

Ich möchte mich auf diesem Wege auch noch einmal für die sehr professionelle Betreuung des Seminars „Optimal Algorithms and Proofs“ letzte Woche bedanken. Ich glaube, das Seminar war sehr gelungen, das persönliche Feedback der Teilnehmer war sehr positiv. Vielen Dank Ihnen und Ihren Kolleginnen für die exzellente administrative Begleitung!

Öffentliche Resonanz im Web 3.3 *Feedback in Social Media*

Mehr und mehr Gäste nutzen die Möglichkeiten des Webs wie Blogs, etc., 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 blogs, Twitter, etc. to share their positive experiences of Dagstuhl with others. Below are some selected excerpts.

Roland van Rijswijk (Radboud University Nijmegen – Nijmegen, NL)

14052 – Ethics in Data Sharing | Dagstuhl Seminar | https://blog.surfnet.nl/?p=3174

Personally, I feel that I learned an incredible amount in a very short time, which I can apply in my work at SURFnet as well as while doing my research. A Dagstuhl seminar is very intense, but worth the effort!

Lance Fortnow (Georgia Institute of Technology – Atlanta, US)

14121 – Computational Complexity of Discrete Problems | Dagstuhl Seminar | https://twitter.com/fortnow/status/445862776112955392

Rite of passage: signing my book for the @dagstuhl library.

Stefan Wagner (Universitat Stuttgart – Stuttgart, DE)

14261 – Software Development Analytics | Dagstuhl Seminar | https://twitter.com/prof_wagnerst/status/480787913366843392

My book at the @dagstuhllibrary. So proud!

Thomas Vogel (Hasso-Plattner-Institut an der Universität Potsdam – Potsdam, DE)

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

Paper “Software Engineering Meets Control Theory”, joint outcome of the GI @dagstuhl seminar 14382 participants, accepted at #SEAMS15

Lance Fortnow (Georgia Institute of Technology – Atlanta, US)

14391 – Algebra in Computational Complexity | Dagstuhl Seminar | http://blog.computationalcomplexity.org/2014/09/dagstuhl-on-algebra-in-computational.html

I learned stuff in the talks, over meals, and even in my room alone at night. […] Looking forward to the next Dagstuhl!

Markus Völter (Ingenieurbüro für Softwaretechnologie – Stuttgart, DE)

14412 – Globalizing Domain-Specific Languages | Dagstuhl Seminar | https://twitter.com/markusvoelter/status/519504474000592896

The library @dagstuhl is very well stocked :-)

Thore Husfeld (The IT University of Copenhagen – Copenhagen, DK)

14451 – Optimality and tight results in parameterized complexity | Dagstuhl Seminar | http://thorehusfeldt.net/2014/11/09/at-first-eppstein-liked-minimum-spanning-trees-4/

I spent a splendid week at Schloss Dagstuhl for Dagstuhl Seminar 14451: Optimality and tight results in parameterized complexity. As usual, this was intellectually extremely stimulating and very exhausting. My head is full of tempting research ideas that I immediately need to pursue.

Fig. 3.3

„My book at the @dagstuhllibrary. So proud!“ Twitter post by 14261 Dagstuhl Seminar participant Stefan Wagner.

Resonanz im Fragebogen 3.4 *Seminar Survey Feedback*

Jeder Seminarteilnehmer erhält von uns einen Fragebogen zur Evaluation der Seminare. Durch dieser anonymen Befragung erhalten wir ebenfalls eine menge positiver Kommentar. Im folgenden zitieren wir hier einige von diesen.

Every seminar participant has the opportunity to fill out a questionnaire about the seminar for evaluation purposes. Below are some excerpts from the many positive comments we received through this anonymous survey in 2014.

14021 – Symmetric Cryptography | Dagstuhl Seminar | http://www.dagstuhl.de/14021

Intense discussions; insightful talks and on a high level.

14022 – Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14022

To me, Dagstuhl is one of the very few occasions that really work regarding scientific collaboration in workshops/conferences.

14022 – Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14022

The smaller size of this Perspectives Workshop was great. It enabled us to focus on this topic, and everyone was given ample opportunity to voice their opinion.

14031 – Randomized Timed and Hybrid Models for Critical Infrastructures | Dagstuhl Seminar | http://www.dagstuhl.de/14031

The organizers did a really good job. It was a pleasure to attend the seminar. I always love to come back to Dagstuhl!

14032 – Planning with epistemic goals | Dagstuhl Seminar | http://www.dagstuhl.de/14032

Superb facilities that were very conducive to research discussion – organization by Dagstuhl staff that made the experience seamless – quality of researchers attending event and thus quality of discussion.

14042 – Do-it-yourself Networking: an Interdisciplinary Approach | Dagstuhl Seminar | http://www.dagstuhl.de/14042

The interdisciplinarity of the seminar was excellent. It was great meeting people from other fields and get new perspectives on things.

14042 – Do-it-yourself Networking: an Interdisciplinary Approach | Dagstuhl Seminar | http://www.dagstuhl.de/14042

A highly enjoyable seminar in a wonderful place. My sincere thanks both to the organisers and to the Dagstuhl organisation and its funders for hosting this event and giving me the opportunity to attend.

14051 – Algorithms for Wireless Communication | Dagstuhl Seminar | http://www.dagstuhl.de/14051

The quality of the presentations was exceptional; I leave with much more complete picture of the area.

14052 – Ethics in Data Sharing | Dagstuhl Seminar | http://www.dagstuhl.de/14052

It was fantastic. Want to do another!

14052 – Ethics in Data Sharing | Dagstuhl Seminar | http://www.dagstuhl.de/14052

The organisation of Dagstuhl itself takes away so much of the practical worries that we have all the time to discuss, this is amazing.

14061 – Statistical Techniques for Translating to Morphologically Rich Languages | Dagstuhl Seminar | http://www.dagstuhl.de/14061

The best aspect is the relaxed and informal interaction with colleagues which leads to discovering mutual research interests and opportunities of collaboration that regular conferences do not facilitate/allow. I could not find a bad aspect really.

14061 – Statistical Techniques for Translating to Morphologically Rich Languages | Dagstuhl Seminar | http://www.dagstuhl.de/14061

Extremely inspiring seminar!!! Thank you 14061 organizers and Dagstuhl staff!!!

14061 – Statistical Techniques for Translating to Morphologically Rich Languages | Dagstuhl Seminar | http://www.dagstuhl.de/14061

Thanks for the good cakes in the afternoon. They make my days. :-)

14062 – The Pacemaker Challenge: Developing Certifiable Medical Devices | Dagstuhl Seminar | http://www.dagstuhl.de/14062

Keep up the great work!

14071 – Graph Modification Problems | Dagstuhl Seminar | http://www.dagstuhl.de/14071

Relaxed friendly environment conclusive to joint research with stimulating talks + plenty of quiet time.

14071 – Graph Modification Problems | Dagstuhl Seminar | http://www.dagstuhl.de/14071

I always enjoy and benefit professionally from coming to Dagstuhl and this time was not an exception.

14072 – New Perspectives in Shape Analysis | Dagstuhl Seminar | http://www.dagstuhl.de/14072

The support from the Dagstuhl office is outstanding.

14081 – Robots Learning from Experiences | Dagstuhl Seminar | http://www.dagstuhl.de/14081

Extremely competent and interesting people to talk with; very stimulating environment; very valuable experience for my PhD students.

14082 – Visualization and Processing of Higher Order Descriptors for Multi-Valued Data | Dagstuhl Seminar | http://www.dagstuhl.de/14082

Out of all the Dagstuhl meetings I have attended, this one provided the most diverse and interesting mix of people from different backgrounds that do not have a comparable opportunity to interact at any other scientific event.

14082 – Visualization and Processing of Higher Order Descriptors for Multi-Valued Data | Dagstuhl Seminar | http://www.dagstuhl.de/14082

I like the long time for discussion and the extremely open atmosphere. People showed their ideas and talked about their plans. They were honestly discussion open questions, failures and got a lot of very helpful suggestions – just great.

14091 – Data Structures and Advanced Models of Computation on Big Data | Dagstuhl Seminar | http://www.dagstuhl.de/14091

Please don't change it. Please.

14092 – Digital Evidence and Forensic Readiness | Dagstuhl Seminar | http://www.dagstuhl.de/14092

Thank you so much for operating Dagstuhl and for providing it to the communities! I truly admire the brains and initiators of the Dagstuhl concept and those who are bring it to life every day again.

14092 – Digital Evidence and Forensic Readiness | Dagstuhl Seminar | http://www.dagstuhl.de/14092

This is the best forum I ever attended! From the bottom of my heart, I thank you! And hope to be invited again!

14092 – Digital Evidence and Forensic Readiness | Dagstuhl Seminar | http://www.dagstuhl.de/14092

Diverse, international, interdisciplinary – just keep it up.

14111 – Combinatorics and Algorithmics of Strings | Dagstuhl Seminar | http://www.dagstuhl.de/14111

The seminar provides a nearly unique opportunity of meeting colleagues beyond the always-busy lifestyle of academia. This format is ideal for free discussions that are often the most fruitful scientifically.

14111 – Combinatorics and Algorithmics of Strings | Dagstuhl Seminar | http://www.dagstuhl.de/14111

I appreciated that the focus of the group is neither too broad nor too narrow – just perfect! The "cross-fertilization" between string algorithms and word combinatorics is very productive and should be continued.

14121 – Computational Complexity of Discrete Problems | Dagstuhl Seminar | http://www.dagstuhl.de/14121

Congratulations for sustained excellence!

14122 – Verification of Cyber-Physical Systems | Dagstuhl Seminar | http://www.dagstuhl.de/14122

Thank you for providing the facility – the format encourages deep & strong technical discussions.

14131 – Computational Models of Cultural Behavior for Human-Agent Interaction | Dagstuhl Seminar | http://www.dagstuhl.de/14131

What an amazing place and opportunity to discuss new ideas, develop collaboration, and define research projects. This should definitely continue. I would love to come back.

14171 – Evaluating Software Verification Systems: Benchmarks and Competitions | Dagstuhl Seminar | http://www.dagstuhl.de/14171

The seminar was absolutely great! I don't know if I want to stay in academia or go to industry after my PhD. This seminar pushed my preferences strongly to direction academia (or industry that is very closely related to academia).

14181 – Multi-agent systems and their role in future energy grids | Dagstuhl Seminar | http://www.dagstuhl.de/14181

The mix between different disciplines was very stimulating; the discussions lively.

14201 – Horn formulas, directed hypergraphs, lattices and closure systems: related formalisms and applications | Dagstuhl Seminar | http://www.dagstuhl.de/14201

The group composition was ideal in terms of people with related, but different research expertise.

14202 – JA4AI – Judgment Aggregation for Artificial Intelligence | Dagstuhl Seminar | http://www.dagstuhl.de/14202

I really appreciated the diverse and expert composition of this seminars group of participant.

14221 – Geometric Modeling | Dagstuhl Seminar | http://www.dagstuhl.de/14221

Organizers did a good job. Staff was extraordinarily accommodating.

14231 – Scientific Visualization | Dagstuhl Seminar | http://www.dagstuhl.de/14231

I have had the pleasure of working with this group for over 15 years. It is an outstanding conference.

14231 – Scientific Visualization | Dagstuhl Seminar | http://www.dagstuhl.de/14231

The group was a nice mix of the "usual suspects" who have been coming to the same meeting for years, and new folks (like me) who have not. I thought that the older group did a great job at integrating the newer into the dialog.

14231 – Scientific Visualization | Dagstuhl Seminar | http://www.dagstuhl.de/14231

It was nice to observe how the community is open to discuss fundamental issues and to jointly address issues that are important for us as a community as a whole (fundamentals, methodological core, curricula) to survive in the ever-changing academic landscape.

14231 – Scientific Visualization | Dagstuhl Seminar | http://www.dagstuhl.de/14231

I believe that these seminars are the best in the world. Two of these have been the most productive professional conferences of my 38-year career (1997, 2003), and this one will be very productive also.

14231 – Scientific Visualization | Dagstuhl Seminar | http://www.dagstuhl.de/14231

Excellent opportunity to brainstorm with like-minded colleagues; discuss controversial area-shaping topics and receive feedback in the most constructive setting ever.

14231 – Scientific Visualization | Dagstuhl Seminar | http://www.dagstuhl.de/14231

This seminar is THE venue where senior researchers in visualization meet.

14232 – Design and Synthesis from Components | Dagstuhl Seminar | http://www.dagstuhl.de/14232

As a junior researcher I felt very much taken seriously by the participants with more experience.

14261 – Software Development Analytics | Dagstuhl Seminar | http://www.dagstuhl.de/14261

Bringing the community together to spend the time to explore an area in depth. I have started a new project with a colleague from another continent based on this workshop.

14261 – Software Development Analytics | Dagstuhl Seminar | http://www.dagstuhl.de/14261

Keep on doing this. I went over my career while I was here and it is amazing how each time Dagstuhl had a mindwarping long term impact on my research in the years after that.

14261 – Software Development Analytics | Dagstuhl Seminar | http://www.dagstuhl.de/14261

Dagstuhl is unique and a great place for sharing and creating ideas – especially in summer.

14281 – Feature Interactions: The Next Generation | Dagstuhl Seminar | http://www.dagstuhl.de/14281

I love coming to Dagstuhl – it's relaxing and inspiring all at the same time. Food quality has continued to improve here over the years – thanks!

14291 – Information-Centric Networking 3 | Dagstuhl Seminar | http://www.dagstuhl.de/14291

Best feature of the seminar was there were about 10 to 15 people who had strong views and had done considerable work in the area. This led to many thought provoking discussions in which the group tried to push past what it had discussed before (sometimes with success, sometimes not, but I felt the group was always working hard). I also thought the outrageous opinion, where everyone got to talk was very important to harvesting ideas from everyone in the workshop.

14302 – Digital Palaeography: New Machines and Old Texts | Dagstuhl Seminar | http://www.dagstuhl.de/14302

This was by far one of the best conferences I had a chance to participate in!

14331 – Querying and Reasoning Under Expressive Constraints | Dagstuhl Seminar | http://www.dagstuhl.de/14331

Mix of inspirational researchers from a broad yet coherent area, and that is never brought together under any normal circumstances.

14341 – Resource-bounded Problem Solving | Dagstuhl Seminar | http://www.dagstuhl.de/14341

Mix – amazingly great.

14342 – Equilibrium Computation | Dagstuhl Seminar | http://www.dagstuhl.de/14342

The organizers did a terrific job. The “lightning talks” were an excellent idea that I haven’t encountered before, and they justified [seminar coordinator] Rahul’s persistence in gathering them. They provided an opportunity for people to touch on work and interests that would not come up later in the lectures, and this in turn led to many interesting conversations.

14361 – Computational Aspects of Fabrication | Dagstuhl Seminar | http://www.dagstuhl.de/14361

The insights from neighboring fields were most useful. We know our peers well and read each others papers. While it is a unique environment to meet peers from the same field, the most useful bits of informations were presented from peers in related fields. Cross collaboration is incredibly important to prevent reinvention of the wheel and accelerating overall progress on pressing research questions. Dagstuhl establishes these cross connections.

14361 – Computational Aspects of Fabrication | Dagstuhl Seminar | http://www.dagstuhl.de/14361

This seminar had an excellent mix of people of different backgrounds, and was very inspiring. It was also interesting that there was some significant overlap in interests with participants of the other seminar taking place at the same time. This also lead to wonderful new contacts and ideas.

14381 – Neural-Symbolic Learning and Reasoning | Dagstuhl Seminar | http://www.dagstuhl.de/14381

Thank you for a truly outstanding scientific experience!

14401 – Privacy and Security in an Age of Surveillance | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14401

I love Dagstuhl. I think it makes a tremendously valuable contribution. This is the second time that a Dagstuhl workshop has helped me launch into a new and very important research area. Thank you.

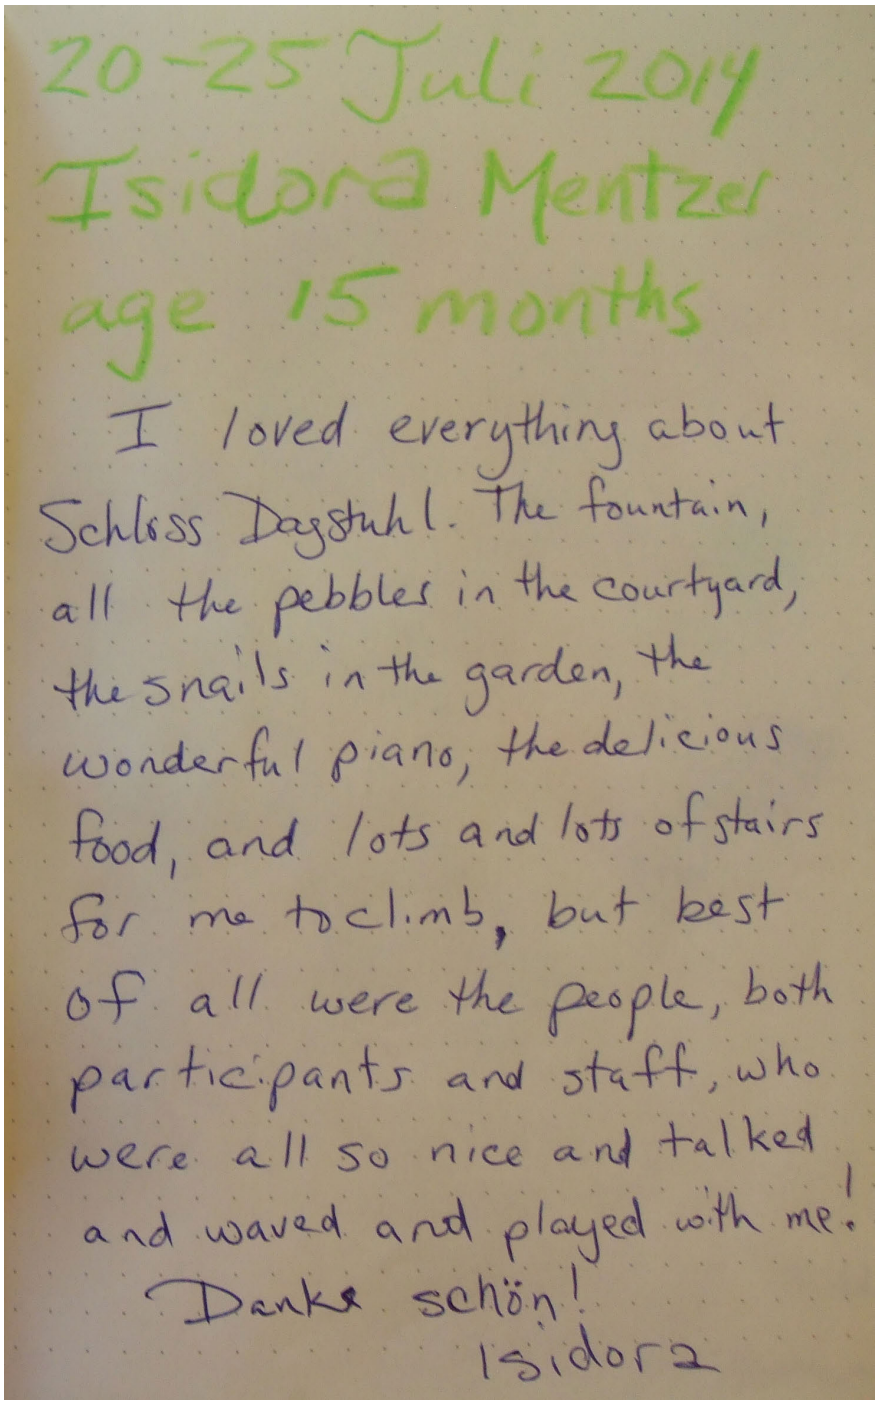

Fig. 3.4

Entry in the Dagstuhl children's guest book by Isidora (15 months), daughter of Dagstuhl Seminar 14301 participant David Smith and partner Cynthia Mentzer, with help from her parents.

„I loved everything about Schloss Dagstuhl. The fountain, all the pebbles in the courtyard, the snails in the garden, the wonderful piano, the delicious food, and lots and lots of stairs for me to climb, but best of all were the people, both participants and staff, who were all so nice and talked and waved and played with me! Danke schön! Isidora and her parents"

14401 – Privacy and Security in an Age of Surveillance | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14401

I thought the balance of fields in the talks, as well as the balance of talks and manifesto work, was excellent, and prevented me from being overwhelmed by too much technical information being delivered at a constant and high rate.

14441 – Modeling, Verification, and Control of Complex Systems for Energy Networks | Dagstuhl Seminar | http://www.dagstuhl.de/14441

The Dagstuhl seminar was a very enjoyable and stimulating week. I have taken away lots of ideas to explore. It was a privilege to participate in this event.

14441 – Modeling, Verification, and Control of Complex Systems for Energy Networks | Dagstuhl Seminar | http://www.dagstuhl.de/14441

Best: insight in another field (power networks) for me as a computer scientist; cooperation possibilities, new inspiring people; resetting from the day-to-day academic hassle at my home university.

14452 – Algorithmic Cheminformatics | Dagstuhl Seminar | http://www.dagstuhl.de/14452

Credit to the organization for the availability of fresh vegetables, salads at lunches, and moderation of sugar and other such junk. I also appreciated the staff's efforts to avoid food wastage. In addition to the practical reasons to do such things, I should note that it creates a good environment to be surrounded by institutional choices that are thoughtful and responsible. Please continue these good priorities.

14452 – Algorithmic Cheminformatics | Dagstuhl Seminar | http://www.dagstuhl.de/14452

While the seminar had a clear focus the background of the participants was very diverse. This resulted in a very stimulating meeting!

14461 – High-performance Graph Algorithms and Applications in Computational Science | Dagstuhl Seminar | http://www.dagstuhl.de/14461

Best: meeting colleagues in a very informal atmosphere with very little external distractions (no family around, no big city distractions, the 5 days period is long enough that there is no time pressure, etc). Even the time to discuss things after talks in and discussion sessions allows things to come up that would not come up in a normal conference.

14471 – Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment | Dagstuhl Seminar | http://www.dagstuhl.de/14471

Best: reunion of some of the finest researchers in the world in the topic being addressed during the workshop, and favorable atmosphere for discussion of existing challenges and novel research directions.

14482 – Automated Planning and Model Checking | Dagstuhl Seminar | http://www.dagstuhl.de/14482

As with previous Dagstuhl seminars, I thought this is a great forum for learning about relevant work in related fields. These seminars play a very useful role and I learned interesting new techniques from other fields and also has the opportunity to connect with researchers in these fields.

Resonanz zur Bibliographiedatenbank dblp

3.5 *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.

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.

Joos Buijs (Eindhoven University of Technology – Eindhoven, NL)
dblp | https://twitter.com/joosbuijs/status/472117256735182848

Thank you @dblp_org for providing an easy to use query API. Writing code to construct a better bibTex file from a messy one was easy!

Andrew McElroy (Codex Labs, LLC, – Nashville, US)
dblp | https://twitter.com/sophrinix/status/505350589556215809

I wish I'd know about this years ago. The DBLP Computer Science Bibliography #compsci #research

Charlotte Truchet (University of Nantes – Nantes, FR)
dblp | https://twitter.com/chtruchet/status/534618982457040896

500 000 distinct users in one month…and still no ads. Thanks a lot @dblp_org !

David Monniaux (VERIMAG – Gières, FR)
dblp | https://twitter.com/dmonniaux/status/559064204305760256

DBLP est le seul index vraiment complet

Moshe Y. Vardi (Rice University – Houston, USA)
dblp | email feedback | http://dblp.dagstuhl.de/

I am a BIG fan of DBLP!

Fig. 3.5
„Beautiful morning on the last day of the seminar here at Schloss Dagstuhl." Twitter post by 14052 Dagstuhl Seminar participant Roland van Rijswijk.

4 Die Seminare in 2014
The 2014 Seminars

Applications, Interdisciplinary Work

- Adjoint Methods in Computational Science, Engineering, and Finance (14371)
- Algorithmic Cheminformatics (14452)
- Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis (14512)
- Computational Aspects of Fabrication (14361)
- Computational Humanities – Bridging the Gap Between Computer Science and Digital Humanities (14301)
- Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing (14022)
- Digital Evidence and Forensic Readiness (14092)
- Digital Palaeography: New Machines and Old Texts (14302)
- Ethics in Data Sharing (14052)
- Exploring Interdisciplinary Grand Challenges in ICT Design to Support Proactive Health and Wellbeing (14272)
- Interaction and Collective Movement Processing (14132)
- Massive Open Online Courses: Current State and Perspectives (14112)
- Multiscale Spatial Computational Systems Biology (14481)
- Randomized Timed and Hybrid Models for Critical Infrastructures (14031)
- Resource-bounded Problem Solving (14341)
- The Pacemaker Challenge: Developing Certifiable Medical Devices (14062)
- Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment (14471)

Artificial Intelligence, Computational Linguistics

- Augmenting Human Memory – Capture and Recall in the Era of Lifelogging (14362)
- Automated Planning and Model Checking (14482)
- Computational Models of Cultural Behavior for Human-Agent Interaction (14131)
- JA4AI – Judgment Aggregation for Artificial Intelligence (14202)
- Neural-Symbolic Learning and Reasoning (14381)
- Planning with Epistemic Goals (14032)
- Preference Learning (14101)
- Robots Learning from Experiences (14081)
- Spatial Reference in the Semantic Web and in Robotics (14142)
- Statistical Techniques for Translating to Morphologically Rich Languages (14061)

Cryptography, Security, Privacy

- Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures (14292)
- Privacy and Security in an Age of Surveillance (14401)
- Socio-Technical Security Metrics (14491)
- Symmetric Cryptography (14021)
- The Synergy Between Programming Languages and Cryptography (14492)

Data Structures, Algorithms, Complexity

- Algebra in Computational Complexity (14391)
- Algorithms for Wireless Communication (14051)
- Analysis of Algorithms Beyond the Worst Case (14372)
- Combinatorics and Algorithmics of Strings (14111)
- Computational Complexity of Discrete Problems (14121)
- Data Structures and Advanced Models of Computation on Big Data (14091)
- Equilibrium Computation (14342)
- Graph Modification Problems (14071)
- Optimal Algorithms and Proofs (14421)
- Optimality and Tight Results in Parameterized Complexity (14451)

Databases, Information Retrieval, Data Mining

- Crowdsourcing and the Semantic Web (14282)
- Querying and Reasoning Under Expressive Constraints (14331)

Distributed Computation, Networks, Architecture, Systems

- Do-it-yourself Networking: an Interdisciplinary Approach (14042)
- High-performance Graph Algorithms and Applications in Computational Science (14461)
- Information-Centric Networking 3 (14291)
- Multi-agent Systems and their Role in Future Energy Grids (14181)
- Resilience in Exascale Computing (14402)

Geometry, Image Processing, Graphics, Visualization

- Geometric Modeling (14221)
- Scientific Visualization (14231)
- Systems and Algorithms for Large-scale Graph Analytics (14462)
- Visualization and Processing of Higher Order Descriptors for Multi-Valued Data (14082)

Software Technology, Programming Languages

- Challenges in Analysing Executables: Scalability, Self-Modifying Code and Synergy (14241)
- Evaluating Software Verification Systems: Benchmarks and Competitions (14171)
- Feature Interactions: The Next Generation (14281)
- Globalizing Domain-Specific Languages (14412)
- Next Generation Static Software Analysis Tools (14352)
- Programming Languages for Big Data (PlanBig) (14511)
- Scripting Languages and Frameworks: Analysis and Verification (14271)
- Software Development Analytics (14261)
- The Future of Refactoring (14211)
- Unifying Product and Software Configuration (14172)

Verification, Logic, Formal Methods, Semantics

- Categorical Methods at the Crossroads (14182)
- Constraints, Optimization and Data (14411)
- Decision Procedures and Abstract Interpretation (14351)
- Design and Synthesis from Components (14232)
- Formal Methods for Coordinating Multi-Agent Systems (14332)
- Horn Formulas, Directed Hypergraphs, Lattices and Closure Systems: Related Formalisms and Applications (14201)
- Modeling, Verification, and Control of Complex Systems for Energy Networks (14441)
- New Perspectives in Shape Analysis (14072)
- Quantitative Models: Expressiveness, Analysis, and New Applications (14041)
- Reachability Problems for Infinite-State Systems (14141)
- Verification of Cyber-Physical Systems (14122)

4.1 Symmetric Cryptography

Organizers: Frederik Armknecht, Helena Handschuh, Tetsu Iwata, and Bart Preneel

Seminar No. 14021

Date: January 5–10, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.1

Participants: Martin R. Albrecht, Elena Andreeva, Frederik Armknecht, Tomer Ashur, Jean-Philippe Aumasson, Steve Babbage, Daniel J. Bernstein, Eli Biham, Alex Biryukov, Céline Blondeau, Andrey Bogdanov, Carlos Cid, Joan Daemen, Itai Dinur, Orr Dunkelman, Henri Gilbert, Jian Guo, Tetsu Iwata, Pascal Junod, Dmitry Khovratovich, Matthias Krause, Tanja Lange, Nils Gregor Leander, Jooyoung Lee, Gaetan Leurent, Eik List, Stefan Lucks, Willi Meier, Florian Mendel, Bart Mennink, Nicky Mouha, Kaisa Nyberg, Kenneth G. Paterson, Thomas Peyrin, Bart Preneel, Christian Rechberger, Greg Rose, Yu Sasaki, Francois-Xavier Standaert, John Steinberger, Gilles Van Assche, Serge Vaudenay, Vesselin Velichkov, Qingju Wang, Jakob Wenzel, Kan Yasuda

Symmetric cryptography deals with the case that both the sender and the receiver of a message are using the same key—the setting for symmetric encryption or authentication—as well as the case where there is no key at all—the setting for cryptographic hash functions. This differentiates symmetric cryptography from it *asymmetric* counterpart, where senders or verifiers use a "public key" and receivers or signers use a corresponding but different "private key." Although asymmetric cryptographic schemes provide in principle more flexibility, but are normally by orders of magnitude less efficient than symmetric cryptographic schemes. Thus, symmetric cryptosystems are the main workhorses of cryptography and highly relevant not only for academia, but also for industrial research, too.

The seminar was the fourth of its kind, the first one took place in 2007, the second in 2009, and the third in 2012. It concentrates on the design and analysis of

- *symmetric primitives* (block and stream ciphers, message authentication codes and hash functions), as well as
- *complex cryptosystems and cryptographic protocols* based on symmetric primitives.

One major topic was authenticated encryption. As already discussed at January 2012 Dagstuhl Seminar on Symmetric Cryptography, there is a demand for encryption schemes that ensure the confidentiality and integrity of data. This eventually led to an open cryptographic competition named CAESAR (Competition for Authenticated Encryption: Security, Applicability, and Robustness)[11] The goal of CAESAR is to identify a portfolio of authenticated ciphers that offer advantages over standard approaches like AES-GCM and (2) are suitable for widespread adoption. To this end cryptographic algorithm designers are invited to submit proposals of authenticated ciphers to CAESAR. All proposals will be made public for evaluation. As the deadline for first round submissions was in March 2014, i.e., only several weeks after the seminar, several groups were actively working on designing and analyzing new proposals for authenticated encryption schemes. Moreover, there was a discussion session that was mainly devoted to current CAESAR submissions. One result was a better understanding of necessary requirements and the current state of these schemes.

Another major topic was the analyis of Even-Mansour encryption schemes. Such schemes generalize common design approaches by reducing these to the composition of simple, idealized components like random permutations. Other topics focused during the discussion session include random number generation and provable security complex cryptosystems.

[11] See http://competitions.cr.yp.to/caesar.html.

Fig. 4.1
Drawing for the Dagstuhl children's guest book by Charlotte (3), daughter of Dagstuhl Seminar 14502 participants Agnes Koschmider and Alexander Paar.

4.2 Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing

Organizers: Peer-Timo Bremer, Bernd Mohr, Valerio Pascucci, and Martin Schulz

Seminar No. 14022

Date: January 5–10, 2014 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.4.1.17

Participants: Abhinav Bhatele, Peer-Timo Bremer, Holger Brunst, Hans-Joachim Bungartz, Remco Chang, Hank Childs, Todd Gamblin, Markus Geimer, Judit Gimenez, Hans Hagen, Daniel A. Keim, Joshua A. Levine, Naoya Maruyama, Bernd Mohr, Christopher Muelder, Klaus Mueller, Matthias S. Müller, Wolfgang E. Nagel, Valerio Pascucci, Ulrich Rüde, Carlos E. Scheidegger, Tobias Schreck, Martin Schulz, Derek Xiaoyu Wang, Felix Wolf

Over the last decades an incredible amount of resources has been devoted to building ever more powerful supercomputers. However, exploiting the full capabilities of these machines is becoming exponentially more difficult with each new generation of hardware. In the systems coming online at this moment, application developers must deal with millions of cores, complex memory hierarchies, heterogeneous system architectures, high-dimensional network topologies as well as a host of other hardware details that may affect the performance of a code. To help understand and optimize the behavior of massively parallel simulations a new subfield of computer science has grown devoted to developing tools and techniques to collect and analyze performance relevant data, such as execution time, operation counts, and memory or network traffic to help application developers pinpoint and ultimately fix performance problems. There now exist a number of standardized tools and APIs to collect a wide range of performance data at the largest scale. However, this success has created a new challenge, as the resulting data is far too large and too complex to be analyzed in a straightforward manner. While there exist some tools for performance analysis and visualization, these are predominately restricted to simple plots of the raw data and rely virtually exclusively on the users to infer connections between measurements and the observed behavior and to draw conclusions. Unfortunately, as the number of cores increases, this approach does not scale. The raw data is typically rather abstract, low-level, and unintuitive and it is difficult to understand within the context of the highly complex interaction of an application with the middle- and system software and the underlying hardware. For this reason, new automatic and more scalable analysis approaches must be developed to allow application developers to intuitively understand the multiple, interdependent effects that their algorithmic choices have on the resulting performance.

Following classical visualization mantra, the natural first step towards automatic analysis is to display an overview of the collected data to provide some insight into general trends. This helps both application developers and performance experts to form new hypotheses on potential causes of and solutions to performance problems. Furthermore, intuitive visualizations are highly effective in conveying the results of any analysis and thus are a valuable tool throughout the entire process. Unfortunately, visualizing performance data has proven challenging as the information is highly abstract, non-spatial, and often categorical. While some early attempts at including more advanced visualizations in performance tools have been proposed, these are rudimentary at best and have not found widespread adoption.

At the same time there exists a vibrant community in the area of information visualization and lately visual analytics that is exclusively aimed at developing techniques to visualize, illustrate, and analyze complex, non-spatial data. In particular, there exists a large body of work on general design principles of visualization tools, color spaces, and user interfaces as well as a wide array of common techniques that tackle a broad range of applications. The Dagstuhl Perspectives Workshop, for the first time, gathered leading experts from both the fields of visualization and performance analysis for joint discussions on existing solutions, open problems, and the potential opportunities for future collaborations.

The week 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. It soon became apparent that there exists a significant overlap in the common tasks and challenges in performance analysis and the abstract problem definitions and concepts common in visualization research. Subsequently, the workshop continued with short talks focusing on various more specific aspects of either existing challenges or potential solutions interspersed with increasingly longer group discussions. Theses extensive, inclu-

sive, and in-depth exchanges ultimately shaped the second half of the workshop and in this form were only made possible through Dagstuhl's unique collaborative and discussion stimulating environment.

Ultimately, the workshop has started a number of collaborations and research projects between previously disparate fields with the potential of significant impact in both areas. Furthermore, the participants distilled the open challenges into three high-level recommendations: First, joined funding for the various open research questions. Second, support to build and foster a new community on the border of visualization and performance analysis. And Third, the need to better integrate the anticipated results into the entire lifecycle of a massively parallel application from design to optimization and production.

4.3 Randomized Timed and Hybrid Models for Critical Infrastructures

Organizers: Erika Ábrahám, Alberto Avritzer, Anne Remke, and William H. Sanders

Seminar No. 14031

Date: January 12–17, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.36

Participants: Alessandro Abate, Erika Ábrahám, Christel Baier, Bernd Becker, Luca Bortolussi, Marc Bouissou, Jeremy T. Bradley, Laura Carnevali, Hermann de Meer, Felicita Di Giandomenico, Martin Fränzle, Hamed Ghasemieh, Marco Gribaudo, Dennis Guck, Lucia Happe, Boudewijn Haverkort, Holger Hermanns, Mauro Iacono, Marijn R. Jongerden, Zbigniew Kalbarczyk, Joost-Pieter Katoen, Anne Koziolek, Peter Langendörfer, Rom Langerak, John Lygeros, Daniel Sadoc Menasche, Albert Molderink, Gethin Norman, Maria Prandini, Anne Remke, William H. Sanders, Sahra Sedigh Sarvestani, Markus Siegle, Gerard J. M. Smit, Oliver Stecklina, Armando Tacchella, Aad van Moorsel, Enrico Vicario, Ralf Wimmer, Verena Wolf, Katinka Wolter

Seminar Description

More and more, our society and economy rely on the well-operation of, often hidden, Information and Communication Technology Infrastructures. These infrastructures play an ever-increasing role in other *Critical Infrastructures*, such as the power grid and water and gas distribution networks. Such systems are highly dynamic and include assets that are essential for the functioning of our society and economy. Users need to be able to place a high level of trust in the operation of such systems, however, uncertainty in the environment, security and physical attacks, and errors in physical devices pose a serious threat to their reliable operation. Hence, it is very important that Critical Infrastructures survive catastrophic events.

Hence, modeling Critical Infrastructures and developing methods to analyze their *safety* and *dependability*, in the presence of failures and disasters is of utmost importance. It is of special interest to analyze, how quickly systems recover to acceptable levels of service after the occurrence of disasters, the so-called *survivability*. However, both failure and repair processes are random and a probability distribution is needed to describe how they evolve over time.

Randomized Timed Models are able to take the dependency of such processes on time into account and powerful techniques exist for their analysis. However, for Critical Infrastructures a modeling formalism is needed that allows describing both discrete and continuous quantities. Examples of discrete quantities are the number of spare parts and the state of sensors, actuators and Information and Communication Technology components, whereas the physical quantities, like the amount of produced energy or the quality of the treated water in terms of temperature and pressure naturally constitute continuous quantities.

Randomized Hybrid Models have been successfully applied to model safety-critical applications. Due to the flexible combination of discrete and continuous state components, Randomized Hybrid Models appear as a natural choice to accurately model Critical Infrastructures. Some formalisms were proposed for the analysis of Randomized Hybrid Models, and an increasing interest and activity can be observed in this field. Still, the industrial application that we are considering is far too large for state-of-the-art approaches; either they are applicable to specific applications only or they do not scale.

Up till now, most modeling in Critical Infrastructures is still fairly "classical" using reliability block diagrams, fault-trees or simplistic stochastic Petri nets. While researchers from the Critical Infrastructures community could benefit from recent advances for Randomized Hybrid Models and their formal analysis, existing algorithms are not yet readily applicable to the special kind of problems arising in Critical Infrastructures.

This clearly shows the need for bringing together experts in the areas of Randomized Timed Models and Randomized Hybrid Models with those from Critical Infrastructures. In the following we describe interesting advances in all three fields and comment on how they can help to bridge the current gap between the fields.

Critical Infrastructures

Critical Infrastructures are in general controlled by SCADA (supervisory control and data analysis) systems, which are potentially vulnerable to attacks and misuse. SCADA systems consist of sensors, actuators, controllers and a human-machine interface through which human operators control the physical process. It is important to correctly capture interdependencies that arise between the SCADA network and the physical network, but also interdependencies between different Critical Infrastructures.

The complex nature of Critical Infrastructures requires a flexible and scalable compositional modeling framework that is able to accommodate different levels of abstraction. At design time, usually not all parameters and not all usage patterns are

known exactly. Also the specific details of vulnerabilities and failures might be unknown, such as the mean time to failure and the impact of a given vulnerability. In such cases it is appropriate to make *stochastic* assumptions about the system and the disaster behavior.

The heterogeneity of typical Critical Infrastructures may require a *combination* of different formalisms and techniques to describe the various components of a system and their dependencies. For example, the combination of continuous and discrete phenomena may need to be captured in the modeling framework, e.g, to model the process automation and the production process which is the essential part of several Critical Infrastructures.

Interactions and dependencies between subsystems of different nature inside a Critical Infrastructure or among cooperating Critical Infrastructures require advanced methods to reconcile different aspects under a common development and assessment framework. *Compositional* modeling can simplify the modeling process and can lead to intuitive formalisms. Furthermore, it enables compositional analysis techniques, which might reduce the complexity of verification and build a challenging topic that requires additional research.

In the seminar we discussed questions like the following:

- Which modeling methods are suitable for which types of Critical Infrastructures?
- Which are the crucial system issues that must be considered when accurately modeling Critical Infrastructures?
- How to distinguish the crucial parameters, thereby keeping the state space of the models as small as possible?

Randomized Timed Models

Randomized Timed Models have been widely used for the modeling and evaluation of, e.g., computer and communication systems. They are in general well understood, suited to model complex systems, and efficient methods and tools exist for their analysis and simulation. Different modeling formalisms differ, e.g., in the model of time (discrete or continuous), in the existence or absence of nondeterminism, or the support of rewards.

Discrete-Time Markov Chains (DTMCs) belong to the most basic probabilistic models, offering a *discretized* model of time in the absence of nondeterminism. Continuous-Time Markov Chains (CTMCs) extend DTMCs by a *continuous* model of time. Several temporal logics were extended to specify relevant properties of Randomized Timed Models, and model checking algorithms were developed to check their validity for the above models. For example, Probabilistic CTL (PCTL) properties for DTMCs can be checked efficiently by solving systems of linear equations. Furthermore, efficient computation algorithms have been developed for model checking Continuous Stochastic Logic (CSL) properties of CTMCs (Baier, Haverkort, Hermanns, Katoen, 2003).

High-level formalisms like General Stochastic Petri Nets (GSPNs) and Stochastic Activity Networks allow to describe complex systems in a more compact way. Their evaluation can be lead back to methods for Markov chains.

Failure and repair processes of Critical Infrastructures often exhibit *nondeterminism*. Markov Decision Processes (MDPs) and Continuous-Time Markov Decision Processes (CTMDPs) extend DTMCs respectively CTMCs with the notion of nondeterminism. These powerful models can be analyzed by determining an optimal scheduler that removes the nondeterminism from the system and allows to apply the model checking approaches for DTMCs and CTMCs. Algorithms exist that compute such optimal schedulers based on solving the underlying optimization problems.

The non-functioning of Critical Infrastructures easily results in huge economic losses. To model the costs of failure and repair, a notion of *reward* can be added to the above models, resulting in so-called Markov Reward Models (MRMs). To specify properties related to rewards, CSL has been extended to Continuous Stochastic Reward Logic (CSRL). Adding rewards to Randomized Timed Models makes the model checking problem very challenging. However, numerical algorithms exist for, e.g., model checking CSRL properties with arbitrary time and reward intervals for CTMCs with rewards. This is extremely useful for Critical Infrastructures, since these algorithms provide a direct and precise method for model checking survivability properties (Cloth, Haverkort, 2005).

There is quite a number of *tools* available for the analysis of the above model types. The most prominent ones are PRISM, MRMC, Möbius, Smart, CADP, or LiQuor. Besides formal verification, there are also simulation-based tools (e.g., APMC, VESTA). Most of these tools were successfully applied to different industrial case studies. However, these formalisms and tools are only partially suited for the model checking of Critical Infrastructures, mainly due to the lack of scalability and modeling power.

Model checking for the above models suffers from the well-known state explosion problem when applied to highly complex and large models of Critical Infrastructures. This problem could be tackled by compositional modeling and verification. However, though the models themselves support compositionality, there are no methods and tools readily available for compositional verification. Moreover, all the above models lack the power to model continuous physical processes, which is an essential part of Critical Infrastructures. Hence, the following section focuses on Randomized Hybrid Models.

In the seminar we discussed questions like the following:

- What are the (dis)advantages of the different modeling formalisms available?
- Which properties of Critical Infrastructures can already be efficiently analyzed with existing techniques?
- What are the requirements for compositional modeling and verification?

Randomized Hybrid Models

When adding continuous behavior to discrete systems, the *hybrid* models become very powerful and in general undecidable. The most popular modeling formalism for hybrid systems are Hybrid Automata. Several analysis techniques were proposed for their reachability analysis, based on, e.g., approximation, hybridization, linearization, the usage of theorem provers, and interval-arithmetic.

Different approaches exist to extend hybrid models with *randomized* behavior. The most important difference between the extensions is *where* randomness is introduced. Timed Automata and Hybrid Automata were extended with *probabilistic discrete jumps* (in the style of DTMCs and MDPs) to Probabilistic Timed Automata respectively Probabilistic Hybrid Automata. In contrast to probabilistic discrete jumps, other formalisms, e.g., Piecewise Deterministic Markov Processes (Davis, 1993), allow *initialized jumps* to take place at *random times* (in the style of CTMCs and CTMDPs).

An orthogonal extension lies in introducing *stochastic differential equations* for modeling perturbations in the dynamic time behavior. When combined with probabilistic discrete jumps, this yields the model of Stochastic Hybrid Systems (Hu, Lygeros, Sastry, 2000). Another possibility considers the combination with

CTMC-style stochastic jumps resulting in Switching Diffusion Processes (Gosh, Araposthatis, Marcus, 1997).

Only some simple classes of these models are decidable; their analysis can be lead back to the analysis of corresponding decidable classes of Hybrid Automata (Sproston, 2000). Despite the undecidability of the above general classes, there are incomplete approaches available for their analysis, based on, e.g., Markov Chain approximation (Prandini, Hu, 2006) or discrete approximation (Koutsoukos, Riley, 2008). Latest work considers CEGAR-style abstraction that allows the application of model checking methods for Hybrid Automata (Zhang, She, Ratschan, Hermanns, Hahn, 2010).

Also the high-level Petri Net models can be extended with hybrid and randomized behavior. Including a notion of time, as in Timed Automata, results in Timed Petri Nets. Hybrid Petri Nets (David, Alla, 2001) are a high-level formalism for general Hybrid Automata. Colored Petri Nets correspond to Piecewise Deterministic Markov Processes (Everdij, Blom, 2009), supporting initialized stochastic jumps. Fluid Stochastic Petri Nets can be seen as a generalization[12] of Piecewise Deterministic Markov Processes, allowing for jumps to take place after a negative exponentially distributed amount of time. Besides the stochastic jumps, these models resolve nondeterminism by introducing discrete probability distributions for concurrently enabled transitions. This way, these models support both a probabilistic choice of jumps and a stochastic randomization of the time point of jumps, making the models extremely expressive and hard to formally analyze. Fluid Stochastic Petri nets can be solved analytically for up to three fluid places. For more general classes, simulation has to be used.

This variety illustrates the emerging interest of the research community in Stochastic and Probabilistic Hybrid Models. Traditionally, academic research focuses stronger on decidable subclasses than on efficient algorithms applicable to more expressive models. However, especially for Critical Infrastructures, models are needed that are able to specify complex continuous dynamics, e.g, in order to study recoverability processes.

For more expressive hybrid models, available analysis methods apply techniques like simulation, dynamic programming, and approximation. The Critical Infrastructures community would strongly benefit from the developments of modern model checking algorithms for models combining randomized and hybrid behavior.

In the seminar we discussed questions like the following:

- What particular hybrid model classes are suitable for Critical Infrastructures?
- How can initialized models be evaluated?
- How can efficient analysis (especially model checking) techniques be adapted for Randomized Hybrid Models?

Achievements of the Research Seminar

This seminar offered a platform to bring together researchers, both from academia and industry, working on *Randomized Timed Models*, *Randomized Hybrid Models* and *Critical Infrastructures*. The program of the seminar was a balanced combination of (i) tutorials and presentations from all three fields to motivate collaboration and to develop a common ground for discussions and (ii) time for collaboration, where actual progress is expected to be made on increasing the efficiency, applicability and application of formal modeling and analysis techniques for Critical Infrastructures.

More specifically, we feel that this seminar helped to improve the development in the given area in the following points:

1. The seminar *increased the interest* for both the academic development and the industrial application of formal methods to Critical Infrastructures and draw attention to open issues. We discussed industrially relevant case studies and specific requirements on modeling formalisms and evaluation techniques in this context.
2. While most of the existing work on Critical Infrastructures focuses on simulation, this seminar aimed at a thorough discussion of the requirements for appropriate formal analysis techniques. We provided an *overview* of the modeling and analysis methods already available in *Randomized Timed and Hybrid Models*, including a thorough discussion of their *suitability* for Critical Infrastructures.
3. We initiated *discussions and cooperations* that advance the state-of-the-art in Critical Infrastructures, regarding both the *development* and the *application* of suitable modeling formalisms and analysis techniques for Critical Infrastructures. We offered a platform to join expertise from different fields, to exchange knowledge about existing methods and applications, to push forward the communication of needs and interests, and to draw attention to challenging research fields and promising applications in the area of Critical Infrastructures.

[12] by skipping the requirement of initialized jumps

Fig. 4.2

Anna-Lena Sauer – Untitled. Part of the Dagstuhl art collection and donated by: David Naumann, Roland Vollmar, Ute Vollmar, Reinhard Wilhelm, participants in Dagstuhl Seminar 03101, and several anonymous donors.

4.4 Planning with Epistemic Goals

Organizers: Thomas Ågotnes, Gerhard Lakemeyer, Benedikt Löwe, and Bernhard Nebel

Seminar No. 14032

Date: January 12–15, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.83

Participants: Thomas Ågotnes, Maduka Attamah, Guillaume Aucher, Christian Becker-Asano, Mikkel Birkegaard Andersen, Thomas Bolander, Jens Claßen, Tiago de Lima, Carmel Domshlak, Hector Geffner, Valentin Goranko, Malte Helmert, Andreas Herzig, Jörg Hoffmann, Martin Holm Jensen, Gerhard Lakemeyer, Jerome Lang, Yongmei Liu, Benedikt Löwe, Robert Mattmüller, Sheila McIlraith, Leora Morgenstern, Bernhard Nebel, Ron Petrick, Gabriele Röger, Francois Schwarzentruber, Sunil Easaw Simon, Hans Van Ditmarsch, Jan van Eijck, Yanjing Wang

Automatic planning is a subarea of Artificial Intelligence that was initiated in the 70s. The main idea was to develop efficient methods to generate action plans, for example for robot missions. The initial attempts were based on first order logic. However, most approaches quickly adapted simpler logics and focused on search techniques. The recent years have brought a huge advance on scalability by employing smart search techniques such as heuristic search, SAT, BDDs, and other techniques. Currently, planning researchers explore widening the scope of planning tasks and to connect back to logic oriented approaches of describing dynamics such as GOLOG. At the same time, planning researchers strive to capture planning settings that are more challenging than the classical setting. For instance, planning under uncertainty and planning taking into account beliefs are current research topics.

The research area of dynamic logics of interaction is part of the larger field of applied and interactive logic: the use of logical methods in order to formalize procedures in social and communication contexts. The systems are typically based on the semantics of modal logic, and often focus on information (ex)change and the dynamics of knowledge and beliefs. Paradigmatic examples are public announcement logic and dynamic epistemic logic. One of the main technical features is the incorporation of agency and events into the modal framework as encapsulated by the notion of product update. Recently, some authors have proposed to use the ideas (or, more generally, the methodology) of dynamic approaches to logic for planning.

Epistemic goals, or more generally, goals that have to be expressed in some intensional language (epistemic, doxastic, deontic, others) have been discussed in several papers in the logic community, but are mostly absent from automatic planning. The development of a research community dealing with these goals in planning will require a close interaction between the two involved communities. The main goal of this workshop was to bring the two communities together and develop a vision of the mid-term goals of such a collaboration. In order to facilitate this, the organizers decided to arrange the workshop around work in four groups: after four tutorial lectures by Gerhard Lakemeyer, Hans van Ditmarsch, Thomas Bolander, and Hector Geffner on Monday, the participants were split up into four groups labelled APPL, BENCH, COMP, and LANG. Tuesday was largely reserved for work in the groups and for preparing the group reports included in this report. Tuesday evening also saw a concert in the *Weisser Saal* with François Schwarzentruber playing the piano and Hans van Ditmarsch playing the cello. The final day had some short presentations and a closing discussion.

For the four themes, the organizers had provided some guiding questions, but left the discussion open for the group participants:

APPL *Applying epistemic planning in the real world.* Theme coordinator: Ron Petrick; group participants: Maduka Attamah, Christian Becker-Asano, Martin Holm Jensen, Benedikt Löwe, Sheila McIlraith, Leora Morgenstern, and François Schwarzentruber. Guiding questions: What are promising applications that will convince the outside world to use epistemic planning? Which areas outside of academia could be interested in epistemic planning? How do we get other academic disciplines (such as roboticists) interested in epistemic planning? Can we come up with a concrete research plan for such an application within the next three years?

BENCH *Establishing benchmarks and concrete goals for epistemic planning.* Theme coordinator: Bernhard Nebel; group participants: Carmel Domshlak, Hector Geffner, Malte Helmert, Andreas Herzig, Jörg Hoffmann, Jérôme Lang, and Hans van Ditmarsch. Guiding questions: Can we come up with standardized problems to measure and compare systems for epistemic planning? Which standardized problems could help to calibrate the expressive power of epistemic planning formalism? What currently unsolved problems will serve as milestones and success criteria for the next three to five years?

COMP *Taming the complexity of epistemic planning*. Theme coordinator: Thomas Bolander; group participants: Gerhard Lakemeyer, Yongmei Liu, Robert Mattmüller, Sunil Simon, Jan van Eijck, and Yanjing Wang. Guiding questions: Which aspects of epistemic planning are responsible for the increase of computational complexity? Are there fragments of epistemic planning that allow for an efficient implementation? Can we devise sufficiently expressive planning formalisms that still have acceptable complexity?

LANG *Finding adequate languages for epistemic planning*. Theme coordinator: Thomas Ågotnes; group participants: Guillaume Aucher, Mikkel Birkegaard Andersen, Jens Claßen, Tiago de Lima, Valentin Goranko, and Gabriele Röger. Guiding questions: Which formalisms are adequate to represent epistemic planning problems? Can we devise languages for epistemic planning that are intuitive to understand and use? Can we extend existing plan definition languages with epistemic features? Are the epistemic logics we have sufficiently expressive to serve as a basis for such planning formalisms?

In the final discussion, the participants discussed the immediate future of the interaction between the two fields. One idea was to edit a special issue of the journal *Annals of Mathematics and Artificial Intelligence*, and the seminar organisers are currently in negotiation with the journal editors about that. Thomas Bolander, Hans van Ditmarsch, Jan van Eijck, and R. Ramanujam are planning a follow-up meeting at the Lorentz Center in Leiden in the spring of 2015, and we hope to reconvene with many of the Dagstuhl participants at that meeting.

4.5 Quantitative Models: Expressiveness, Analysis, and New Applications

Organizers: Manfred Droste, Paul Gastin, Kim Guldstrand Larsen, and Axel Legay

Seminar No. 14041

Date: January 19–24, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.104

Participants: Giorgio Bacci, Giovanni Bacci, Benedikt Bollig, Miroslav Ćirić, Thomas Colcombet, Julie De Pril, Laurent Doyen, Manfred Droste, Zoltán Ésik, Javier Esparza, Ulrich Fahrenberg, Emmanuel Filiot, Paul Gastin, Daniel Gebler, Cyrille Jegourel, Heinz Koeppl, Jan Křetínský, Antonín Kučera, Dietrich Kuske, Kai Lampka, Kim Guldstrand Larsen, Axel Legay, Björn Lisper, Sylvain Lombardy, Jan Madsen, Nicolas Markey, Benjamin Monmege, Vitaly Perevoshchikov, Tatjana Petrov, Karin Quaas, Christian M. Reidys, Martin Riedl, Cristian Riveros, Jacques Sakarovitch, Sean Sedwards, Gerard J. M. Smit, Jeremy Sproston, Bart Theelen, Thomas Weidner, Rafael Wisniewski

Quantitative models and quantitative analysis in Computer Science is receiving increased attention in order to meet the challenges from application areas such as Cyber Physical Systems. What is aimed at is a revision of the foundation of Computer Science where Boolean models and analyses are replaced by quantitative models and analyses in order that more detailed and practically useful answers can be provided. Recently, a large number of new models, toolsets, and new application domains have emerged. The theory of weighted automata has also developed, introducing extensions of the models which are motivated by the quantitative analysis of systems.

The first objective of the seminar was to bring the quantitative model checking and weighted automata communities together with the goal of discussing the latest developments in those areas. The second objective of this workshop was to go one major step further. In fact, it has been recently observed an increasing usage (and demand) of quantitative models in a wide range of new application domains. This includes, e.g., systems biology and energy grid. However, these different communities are often not aware of each other. This seminar had the major objective to put those various communities in contact with the hope of creating fruitful long term collaborations.

Quantitative model checking covers extended automata-based models that permit to reason on quantities. The model of timed automata introduced by Alur and Dill in 1989 has by now established itself as a universal formalism for describing real-time systems. The notion of zone has led to a number of tools – e.g. BIP, Kronos, UPPAAL – which support efficient analysis (reachability and model checking) of timed automata. Later the more expressive formalism of hybrid automata was introduced and popularized by Henzinger et al and the introduction of the tool HyTech provided a semi-decision algorithm for analyzing so-called linear hybrid systems. Whereas in timed automata the continuous part of a model is restricted to be clocks (which always evolve with rate 1), linear hybrid automata allow more general continuous variables with evolution rates in arbitrary intervals. The notion of priced (or weighted) timed automata was introduced independently by Alur et al and Larsen et al in 2001, with the surprising result that cost optimal reachability is decidable. Since these initial results, efficient tools were developed and a number of more challenging questions have been considered including multi-priced timed automata, optimal infinite scheduling (both with respect to mean pay-off and discounting), priced timed games and model checking for priced timed automata.

Driven by new needs in areas such as cyber physical systems, a series of recent work have tried to combine real-time with stochastic aspects, leading to new models such as timed stochastic automata. One of the main objectives of the seminar was to study those new models and put them in perspective with similar results in weighted automata. The new notion of energy automata (Larsen, Markey, Bouyer, ...) that extends price timed automata and permits to reason on energy problems was also discussed and put in perspective with similar work done at the weighted automata level.

Weighted automata on finite words were already investigated in seminal work of Schützenberger (1961) and Chomsky and Schützenberger (1963). They consist of classical finite automata in which the transitions carry weights which may model, e.g., the cost, the consumption of resources, or the reliability or probability of the successful execution of the transitions. This concept soon developed a flourishing theory. Recently, motivated by practical examples of energy consumption, new quantitative automata models have been introduced and investigated in which the weights of finite or infinite paths are computed e.g. by the average weights or by the accumulation points of the average weights of their transitions. Colcombet (2009) studied regular cost functions which permit a quantitative extension of classical equivalence results relating automata, expressions, algebraic rec-

ognizability, and variants of monadic second-order logic. Gastin et al (2010) introduced weighted pebble automata in order to capture the expressive power of weighted extensions of Xpath for XML documents, or temporal logics for linear behaviors. All these concepts provide totally new models for which weighted automata-theoretic methods can often be applied successfully. It was very profitable therefore to bring these different communities together.

Another main theme of the seminar was to create interaction with researchers working in areas where the theoretical models and techniques may have potential applications. In systems biology, the challenge is not only to find mathematical models, but also to define new efficient quantitative analysis techniques capable of coping with very large size complex systems. Two promising applications are 1) using SMC-based techniques to monitor complex properties that cannot be expressed in classical temporal logic (e.g., oscillation properties), and 2) using interface theories as a formal characterization of phenomena in the area of synthetic biology. As another application area, the challenge of smart electricity grids is to balance the behavior of all participants (suppliers and consumers) to improve efficiency and stability. Again, quantitative models such as energy automata and analysis are emerging as potential key techniques.

In the seminar, 40 researchers from 13 countries discussed their recent research results and developments for quantitative models and their analysis. Five survey lectures, including two lectures covering the application domains, and 32 talks were organized in eight sessions with centralized themes. From the beginning, all lectures and talks raised questions of members from the other fields, and lively discussions followed. In particular, the surveys presented the fields of weighted automata, formal model checking and simulation methods adopted by industry, programmable single-cell biocomputers, models for smart grid balancing, and asymptotic analysis of weighted automata. The lectures and talks dealt with, e.g., quantitative logics and their semantics, expressiveness of models including quantitative measures for infinite behavior (like discounting, mean payoff, long-run averages), and statistical model checking of stochastic hybrid systems, to name only a few topics.

There are a number of open problems concerning the interplay between these fields. For instance, there are many interesting open questions about the connection between energy automata, energy functions and weighted automata, on weighted specification languages used in more algebraic settings, on energy games, and on the combination of real-time and probabilism. The interplay between priced timed automata and weighted automata also demands further investigation. Due to these open challenges, several researchers decided to meet again later in the year, e.g. during the international workshop in Leipzig on “Weighted Automata: Theory and Applications (WATA 2014)”.

During the seminar, there was very much interaction between the participants. In particular, the seminar was successful in attracting academic researchers with contacts to industry; this was felt very positive and should definitely be continued. Generally, it was expressed that a future research collaboration between the different present groups should be highly fruitful and would therefore be very desirable. A Dagstuhl seminar would provide an ideal and unique opportunity for this. The successful collaboration in the present seminar was felt to be due in particular to the superb facilities and excellent organization provided by the Dagstuhl center and its team.

4.6 Do-it-yourself Networking: an Interdisciplinary Approach

Organizers: Panayotis Antoniadis, Jörg Ott, and Andrea Passarella

Seminar No. 14042

Date: January 19–22, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.125

Participants: Panayotis Antoniadis, Ileana Apostol, N. Asokan, Jonathan Baldwin, Christian Becker, Jon Crowcroft, Fiorella De Cindio, Paul Dourish, Kevin R. Fall, Marcus Foth, Mark Gaved, Per Gunningberg, Ahmed Helmy, Paul Houghton, Katrina Jungnickel, Teemu Kärkkäinen, Jussi Kangasharju, Gunnar Karlsson, Anders Lindgren, Marcin Nagy, Christian Nold, Jörg Ott, Andrea Passarella, Dan Phiffer, Alison Powell, Amalia Sabiescu, Douglas Schuler, Irina Shklovski, Michael Smyth, Ersin Uzun, Volker Wulf

The key objective of the seminar was to bring together a diverse group of researchers and practitioners to reflect on technological and social issues related to the use of local wireless networks that operate outside the public Internet. We managed to bring together a quite balanced group of 32 people with expertise in the design and implementation of wireless ad hoc networks of various types, human-computer interaction, community informatics, urban interaction design, ethnography, media studies, arts and design.

Interdisciplinary interactions took place successfully around specific application areas for which the use of do-it-yourself networks is meaningful. More specifically, we explored the use of such networks for supporting the creation of transient communities of different size and duration, political activism, and similarity matching. In addition, an in depth exploration of the concept of failure provided a useful framework for addressing various challenges in bridging the gap between theory and practice, scientific and social objectives.

Our main finding was that there are certain assumptions that need to be carefully understood and important requirements that need to be fulfilled in order for DIY networking to become a feasible, and desirable, option for shaping the hybrid space of contemporary cities. That calls for a closer collaboration between experts from different fields and disciplines. For this, the most important achievement of our seminar was the balanced and productive interactions between engineers and social scientists around a concrete topic, and the general feeling that a new interdisciplinary community around the topic of DIY networking is meaningful and a goal worth pursuing. Indeed, concrete plans for facilitating the formation and expansion of such a community through online communication and face-to-face meetings, research visits, and common projects between participants that met in Dagstuhl for the first time are already under way.

When things get so big, I don't trust them at all
You want some control – you've got to keep it small
D.I.Y. D.I.Y. D.I.Y. D.I.Y.
— Peter Gabriel

Fig. 4.3
Drawing for the Dagstuhl children's guest book by Konrad (3), son of Dagstuhl Seminar 14502 participants Agnes Koschmider and Alexander Paar.

4.7 Algorithms for Wireless Communication

Organizers: Guy Even, Magnús M. Halldórsson, Yvonne Anne Pignolet, and Christian Scheideler

Seminar No. 14051

Date: January 26–31, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.152

Participants: Eyjólfur Ingi Asgeirsson, Marijke Bodlaender, Johannes Dams, Michael Dinitz, Yuval Emek, Guy Even, Sándor Fekete, Fabian Fuchs, Jie Gao, Seth Gilbert, Martin Haenggi, Bernhard Haeupler, Magnus M. Halldorsson, Martin Hoefer, Stephan Holzer, Qiang-Sheng Hua, Thomas Janson, Tomasz Jurdzinski, Holger Karl, Matthew J. Katz, Thomas Kesselheim, Christian Konrad, Bhaskar Krishnamachari, Fabian Daniel Kuhn, Bodo Manthey, Friedhelm Meyer auf der Heide, Calvin Newport, Merav Parter, Marina Petrova, Yvonne-Anne Pignolet, Dror Rawitz, Kay Römer, Adi Rosén, Alexander Russell, Christian Scheideler, Christian Schindelhauer, Ravi Sundaram, Takeshi Tokuyama, Roger Wattenhofer

The last decades have seen an ever growing interest in wireless communication networks and their applications. Wireless networks pose many algorithmic challenges for various reasons: Realistic wireless signal propagation and interference models are very complex and therefore hard to use in rigorous algorithmic research, and this is further complicated by emerging technologies such as MIMO (multiple-input and multiple-output). Also, reasonable models for the dynamics and mobility in these networks can be quite complex and are not yet well-understood. Furthermore, standard complexity measures such as time and space are not sufficient any more as energy consumption is also a critical aspect that cannot be neglected. Many protocols for wireless networks have already been proposed by the research community, but most of them have only been studied in simulations or analyzed using rather simple models. So there is doubt whether any of these protocols would actually work in practice.

The purpose of this Dagstuhl seminar was to bring together computer scientists of different backgrounds to review and discuss models and algorithmic approaches in order to obtain a better understanding of the capabilities and limitations of modern wireless networks and to come up with more realistic models and approaches for future research on wireless networks that may then be investigated in joint research projects. The mix of the participating people resulted in fruitful discussions and interesting information exchange. The structure of the seminar took advantage of these different backgrounds by focusing on themed talks and open discussions.

The program included an eclectic mix of algorithmic and systems perspectives, modeling issues and emerging networking techniques, and explorations of the limits and possibilities of fundamental problems.

Discussions of models ranged from simple graph-based communication and interference models, to stochastic models, adversarial interruptions and jamming, dynamic networks and uncertainty formulations, and variations and extensions of signal-strength models.

Presentations from the systems perspective included managing environmental factors affecting measurements, robust predictions of channel capacities, efficiency of backpressure routing, issues in emerging heterogeneous radio environmental contexts, and robots controlled via wireless communication.

New dimensions at different networking layers included MIMO, network coding, interference cancellation, directional antennas and cognitive radio networks.

Finally, new results were presented on various related classic problems including broadcast, local broadcast, game theory, coding, routing, positioning, and connectivity.

4

4.8 Ethics in Data Sharing

Organizers: Julie Cohen, Sven Dietrich, Aiko Pras, Lenore D. Zuck, and Mireille Hildebrand
Seminar No. 14052

Date: January 26–31, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.1.170

Participants: Jon Callas, Georg Carle, Julie E. Cohen, Sven Dietrich, Ronald Leenes, Aiko Pras, Volker Roth, Peter Y. A. Ryan, Jürgen Schönwälder, Darren Shou, Anna Sperotto, Radu State, Burkhard Stiller, Jeroen van der Ham, Roland van Rijswijk-Deij, Aimee van Wynsberghe, Da-Wei Wang, Sam Weber, Lenore D. Zuck

ACM's ethical guidelines (as well as IEEE's) are almost two decades old. The most relevant points to data sharing it makes are "Avoid harm to others" and "Respect the privacy of others." The consequences of not complying with the code are "Treat violations of this code as inconsistent with membership in the ACM" while "Adherence of professionals to a code of ethics is largely a voluntary matter."

In fact, in the current legal system, ethical behavior "doesn't pay." Such guidelines are insufficient for the numerous professionals working for corporations where privacy policies are dictated more by a company than by its employees. Nowadays, we have little control who receives our Personally Identifiable Information (PII), what PII they receive, where collected PII is transferred to, and what is the source of (mis?)information others have on us. This is especially alarming with the rapid progress of data mining, the constant discovery of flaws in data anonymization/sanitization techniques, and the vast amount of electronic data that exists. It is beyond the ability of a layperson to understand the privacy policy of organizations and their consequences on the individual.

The situation is even more serious when data is shared and disseminated among different countries that naturally have different ethical codes and policies for dealing with privacy issues concerning data sharing. Data transfer has no borders, hence, neither does data sharing, which renders ethical data sharing all the more challenging.

However, the recent EU proposals to update the legal framework of the Fair Information Principles, precisely with an eye to the emergence of hyperconnectivity and ubiquitous data analytics, has introduced the notion of Data Protection by Design. This may provide strong incentives to introduce purpose binding, informed consent, minimal disclosure and profile transparency into the design of the relevant computing systems.

The seminar brought in researchers from all disciplines that involve data sharing across borders with ethical implications. The main focus was on Computer System Security data, with consideration for Electronic Medical Records. We derived a basic model for data sharing, and came up with some suggestions of code of ethics for computer professionals (including researchers) that will elaborate on existing codes in terms of data sharing.

4.9 Statistical Techniques for Translating to Morphologically Rich Languages

Organizers: Alexander Fraser, Kevin Knight, Philipp Koehn, Helmut Schmid, and Hans Uszkoreit

Seminar No. 14061

Date: February 2–7, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.1

Participants: Arianna Bisazza, Fabienne Braune, Fabienne Cap, Marine Carpuat, David Chiang, Ann Clifton, Hal Daumé III, Gideon Maillette de Buy Wenniger, Chris Dyer, Andreas Eisele, Richard Farkas, Marcello Federico, Mark Fishel, Anette Frank, Alexander M. Fraser, Spence Green, Nizar Habash, Jan Hajič, Katrin Kirchhoff, Kevin Knight, Philipp Koehn, Jonas Kuhn, Alon Lavie, Krister Linden, Andreas Maletti, Maria Nadejde, Preslav Nakov, Hermann Ney, Joakim Nivre, Slav Petrov, Maja Popovic, Anita Ramm, Stefan Riezler, Hassan Sajjad, Helmut Schmid, Hinrich Schütze, Khalil Sima'an, Sara Stymne, Jörg Tiedemann, Kristina Toutanova, Hans Uszkoreit, Josef van Genabith, Sami Virpioja, Stephan Vogel, Martin Volk, Bonnie Webber, Marion Weller, Phil Williams, Shuly Wintner, Dekai Wu, François Yvon

The seminar on Statistical Techniques for Translating to Morphologically Rich Languages allowed disparate communities working on problems related to morphologically rich languages to meet to discuss an important research problem, translation to morphologically rich languages. While statistical techniques for machine translation have made significant progress in the last 20 years, results for translating to morphologically rich languages are still mixed versus previous generation rule-based systems, so this is a critical and timely topic. Current research in statistical techniques for translating to morphologically rich languages varies greatly in the amount of linguistic knowledge used and the form of this linguistic knowledge. This varies most strongly by target language, for instance the resources currently used for translating to Czech are very different from those used for translating to German. The seminar met a pressing need to discuss the issues involved in these translation tasks in a more broad venue than the ACL Workshops on Machine Translation, which are primarily attended by statistical machine translation researchers.

Important background for the discussion was the recent realization that more linguistically sophisticated methods are required to solve many of the problems of translating to morphologically rich languages. Therefore it was critically important that SMT[13] researchers be able to interact with experts in statistical parsing and morphology who work with morphologically rich languages to discuss what sort of representations of linguistic features are appropriate and which linguistic features can be accurately determined by state of the art disambiguation techniques. This was an important step in creating a new community crossing these research areas. Additionally, a few experts in structured prediction were invited. The discussions took advantage of their insight in how to jointly model some of these phenomena, rather than combining separate tools in ad-hoc pipelines as is currently done. The overall discussion was driven by the following questions:

- Which linguistic features (from syntax, morphology and other areas such as coreference resolution) need to be modeled in SMT?
- Which statistical models and tools should be used to annotate linguistic features on training data useful for SMT modeling?
- How can we integrate these features into existing SMT models?
- Which structured prediction techniques and types of features are appropriate for training the extended models and determining the best output translations?
- What data sets should be used to allow a common test bed for evaluation?
- How should evaluation be conducted, given the poor results of current automatic evaluation metrics on morphologically rich languages?

The Dagstuhl seminar on Statistical Techniques for Translating to Morphologically Rich Languages addressed these questions by allowing four different communities to meet together: statistical machine translation, statistical parsing, morphology and structured prediction.

Outcome in brief. The Dagstuhl seminar on Statistical Techniques for Translating to Morphologically Rich Languages was a great success. The discussions held will play an important role in allowing researchers to significantly advance the state-of-the-art. In particular, strong and weak points in current research approaches were identified and proposals to address the weak points were made. In addition, the seminar acted as a valuable venue for more junior researchers to spend more time

[13] SMT – Statistical Machine Translation

talking with senior researchers than is possible in a conference setting. Finally, several new community building ideas were discussed, including a DFG proposal connecting all of the major sites for statistical machine translation research in Germany, see below.

Invited Talks. We begin the detailed discussion with a brief idea about the three invited keynote talks (as well as the introductory overview and motivational talk). All of these talks were very well received, with several seminar participants commenting that they learned a significant amount by being able to see a synthesis of the problems, current approaches and possible future approaches to translating to morphologically rich languages. The three keynote talks were:

- Philipp Koehn of the University of Edinburgh presented a general discussion of dealing with the phenomena of morphologically rich languages in translation.
- Kristina Toutanova of Microsoft Research presented a detailed overview of the state-of-the-art in statistical machine translation research related to morphologically rich languages in translation.
- Kevin Knight of the University of Southern California presented a vision of the future, where the field could go, in terms of both better modelling of morphologically rich languages, and the use of more language independent structure (at the semantic level) in translation.

After this, people interested in leading a discussion group held talks.

Discussion Groups. There were initially nine proposed topics for discussion groups (note that these are listed as topic-focused talks subsequently in the report):

- Nivre/Petrov: Parallel dependency treebanks and linguistic resources
- Tiedemann: The use of synthetic training data and pivot languages to overcome data sparseness
- Kirchhoff: Language modeling
- Dyer: Modeling morphemes vs. modeling words and smoothing with morphemes
- Habash: Arabic morphology and deep morphology representation for MT
- Williams/Koehn: Syntactic SMT for morphologically rich languages
- Knight: Semantics
- Webber: Discourse/aspects of semantics
- Bojar/Hajič: Generating morphology for SMT

Following this all participants emailed the organizers with their discussion group preferences. In the end, all but two participants were assigned to their first preference. We eliminated two groups (on synthetic training data and generating morphology), and their proposers joined other groups.

Following initial group presentations by some groups on Wednesday morning, three groups dissolved and several decided to continue. The three new groups that were proposed were:

- Virpiojia/Dyer: Unsupervised morphology for statistical machine translation
- Wu/Lavie: Evaluation of machine translation output
- Nivre/Knight: Universal Annotation and Abstract Meaning Representation

Highlights of what was accomplished by the discussion groups were:

- Dyer and Virpiojia and groups looked at morphologically aware translation models which use morphology to cover the long-tail without requiring morphological modelling of very frequent tokens, and looked at the state-of-the-art in unsupervised modeling.
- Kirchhoff and her group carried out a detailed survey of the state-of-the-art for language modeling of morphologically rich languages and documented this on the Wiki.
- Nivre and his two groups (one co-led with Petrov) defined a new proposed annotation standard for working on two levels (surface forms and lemmas, including multi-word-entities and decomposed compounds).
- Habash and his group carried out a literature review of attempts to deal with Arabic morphology in translation, discussing the strengths and weaknesses of the approaches, and identifying a new direction for future work.
- Williams, Koehn and group looked at the application of unification to modelling agreement in multiple languages.
- Knight and his two groups worked on general applications of semantically-aware processing to morphologically rich languages and on identifying areas where the Abstract Meaning Representation could be applied to this problem.
- Webber and group created a list of resources and research papers on applying discourse modeling to statistical machine translation and looked at machine translation output to find errors caused by broken discourse constraints.
- Wu, Lavie and group discussed and documented the different levels of linguistic analysis required for high quality automatic evaluation when the target language is morphologically rich.

See the individual abstracts for more information and further details.

Other activities. In addition to the formal work carried out in the talks and discussion groups, Dagstuhl offered an intimate environment strongly encouraging networking and discussion. The meal system of Dagstuhl, with random assignment of people to tables, is an excellent idea and was particularly useful for the more junior participants who did not know many of the senior researchers attending (several people mentioned informally that this was the best experience of this sort they have had). The informal evening activities centering around social gatherings and the music room were also very well attended and a variety of interesting discussions took place. The excursion to Trier was a welcome mid-week break and provided another networking opportunity, as well as being highly interesting for the vast majority of participants who had not previously visited a city with a similar historical background.

The seminar was unusual for Dagstuhl itself in that very few of the participants had participated in a Dagstuhl seminar previously. Due to the strongly positive reaction we anticipate that other research areas within Natural Language Processing will apply for Dagstuhl seminars.

We would like to take the opportunity here to thank Dagstuhl for the wonderful logistic support and for providing such a stimulating environment for our work.

Communities represented in more detail. The seminar was a success in terms of the strong participation of women and a good geographical distribution (although Asia could have been somewhat more strongly represented). Our only strong area of concern was that of the numerous participants from companies invited, only two attended (Kristina Toutanova of Microsoft Research and Slav Petrov of Google, who gave one of the keynotes and co-led a discussion group respectively). Nevertheless the networking opportunities were excellent and many participants informally told us that this was an excellent meeting which they expected to have a strong impact on their research.

One characteristic of the proposal which was successfully

carried out was a meeting of four different communities: statistical machine translation, statistical parsing, morphology and structured prediction. In particular, we felt that the interaction between the statistical machine translation researchers and the researchers working on statistical parsing and morphology was highly productive and will likely lead to new techniques of analyzing morphologically rich languages which will be more useful in translation research than the current approaches. We believe that the Dagstuhl seminar has been unique in terms of providing the opportunity for these communities to meet together for five days and understand each others' perspective on research.

Conclusion and Impact. In conclusion, we believe the Dagstuhl seminar has met the goals we set out for it, in terms of providing a forum for discussion of the current problems with the state-of-the-art and allowing a focusing of research effort which was not previously present in the research community.

As we previously mentioned, in addition to the less quantifiable aspects in terms of networking and connections made, there were several prominent concrete outcomes of the Dagstuhl seminar. The new annotation standard suggested by the two Universal Annotation groups led by Nivre, Petrov and Knight is one strong outcome which will change the basic tools that the statistical machine translation community will have available. The Kirchhoff group is working on a position paper that will help to refocus effort on language modeling for morphologically rich languages, which will have an impact not only on machine translation research but also research on speech recognition and other research areas.

Five of the six most prominent researchers in machine translation in Germany were able to attend the Dagstuhl seminar, and while there have decided to launch a new research program in translating spoken language in an educational context, with a particular focus on translation to German (a morphologically rich language), by submitting a Paketantrag to the DFG. The work will be carried out with a view toward creating a DFG Schwerpunktprogramm focusing on Natural Language Processing for German after the successful completion of the work in the Paketantrag. The researchers are Fraser, van Genabith, Ney, Riezler, Uszkoreit, and they are joined by Alex Waibel (who was invited to the seminar but unable to attend due to scheduling conflicts). This new funding effort would not have been possible without the possibility to meet at Dagstuhl several times to find common ground and determine an overall strategy.

In short, we were very happy with the discussions, work and impact of the Dagstuhl seminar on translation to morphologically rich languages. We plan to apply to hold a second meeting at Dagstuhl in the summer of 2016 on the same topic.

Finally, we would like to once again thank the staff of Dagstuhl for facilitating these unique scientific discussions which we are confident will have a strong impact on future research on the important problem of statistical techniques for translation to morphologically rich languages.

Fig. 4.4
Fabian Treiber – Apart from Mauve. Part of the Dagstuhl art collection and donated by: Rogerio de Lemos, Jörg Ott, Roland & Ute Vollmar, Reinhard Wilhelm, participants in Dagstuhl Seminar 14052, and various anonymous donors.

4.10 The Pacemaker Challenge: Developing Certifiable Medical Devices

Organizers: Dominique Méry, Bernhard Schätz, and Alan Wassyng

Seminar No. 14062

Date: February 2–7, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.17

Participants: Andrew Butterfield, Martin Daumer, John S. Fitzgerald, Michaela Huhn, Zhihao Jiang, Soeren Kemmann, Günther Klebes, John Komp, Peter Gorm Larsen, Brian Larson, Mark Lawford, Zhiming Liu, Dominique Méry, Zaur Molotnikov, Pieter J. Mosterman, Marcel Oliveira, Christian Prehofer, Florian Prester, Francesca Saglietti, Bernhard Schätz, Neeraj Kumar Singh, Markus Völter, Alan Wassyng, Reinhard Wilhelm

Pacemakers are typical examples of those medical devices, like insulin pumps, that help save lives when they operate correctly and safely, but may cause grievous harm when they fail. State-of-the art safety standards like IEC 61508 highly recommend (semi-)formal methods for the specification, design, and development of those devices. The Pacemaker Formal Methods Challenge, the first challenge issued by the North American Software Certification Consortium, is hosted by the Software Quality Research Lab at McMaster University, Canada. The challenge is based on a pacemaker specification offered by Boston Scientific, and is part of the verification Grand Challenges which is an international, long-term research programme that seeks to create a substantial and useful body of code that has been verified to the highest standards of rigour and accuracy. The Pacemaker case-study attracted substantial participation during different events in the research community such as workshops at FM2008, FM2009, FHIES 2011, FHIES 2012 and the student competition at ICSE2009 (SCORE). Currently there are more than 10 world-class research institutes and universities that take part in the challenge, and are using different approaches. Today, there is a wide range of approaches in the formal methods community to specify and develop high integrity systems. Many of these formal approaches do not work well on industrial level applications, and so the state of the practice is remarkably deficient, even in the case of systems that require certification according to the highest safety levels. The purpose of this five days seminar was to bring together researchers, regulators, as well as practitioners in the medical field to discuss and compare different approaches for the development of certifiable medical software, and further the state of practice. Listed below are research topics related to development of medical software to be covered in the seminar:

- Certification: How can formal methods help in the process of certification of embedded medical software? What standards are in current use and in what measure do they cover model based development? How do we address safety, security and privacy now that these implantable devices are equipped with Wi-Fi, Bluetooth and other wireless networking technologies? How do unspecified environmental assumptions affect the final product?
- Model-based Development: How can established methods for model based development help the building of implantable medical devices? What kind of models (e.g. controlled biological process, hardware platform, safety function) are needed for designing and certifying safety critical medical systems?
- Medical-domain specific aspects: What are the most important specific non-functional aspects that need to be considered while developing implantable medical devices? How can biological and medical aspects be integrated in the development process?
- Tooling: What is the current state of the art and practice concerning tools for formal specification that would be useful in the medical device domain?
- Pragmatics: What is the fitness of different methods for transfer into practice? What do we need to do to ensure that the regulators and workforce are adequately informed of methods and tools that are useful/indispensable in this domain?

As major results of this Dagstuhl Seminar, two publications are targeted at all three relevant sectors researchers, regulators and manufacturers.

The first outcome is a comparison of the different approaches to the Pacemaker Challenge, to be available as a Dagstuhl publication. To achieve such a comparison, the organizers have prepared a catalogue of criteria according to which the approaches are compared. This catalogue was available in advance of

the seminar, so presenters can provide a rationale for their classification according to the catalogue, and participants can discuss those classifications.

As a further, more formal result, a joint publication most preferably in the form of a book on the use of rigorous methods for the development of software-intensive medical devices with the pacemaker as a common example will be produced, with the organizers and editors, and all invited research groups as co-authors. Commitment to the participation in this publication will be made a prerequisite for participation in the seminar for members of the research groups having participated in the challenge.

4.11 Graph Modification Problems

Organizers: Hans L. Bodlaender, Pinar Heggernes, and Daniel Lokshtanov

Seminar No. 14071

Date: February 9–14, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.38

Participants: Isolde Adler, Rémy Belmonte, Hans L. Bodlaender, Andreas Brandstädt, Yixin Cao, Derek G. Corneil, Marek Cygan, Tinaz Ekim, Henning Fernau, Fedor V. Fomin, Archontia C. Giannopoulou, Petr A. Golovach, Martin Charles Golumbic, Alexander Grigoriev, Pinar Heggernes, Pim van 't Hof, Bart M. P. Jansen, Mamadou Moustapha Kanté, Eun Jung Kim, Dieter Kratsch, Stefan Kratsch, O-joung Kwon, Van Bang Le, Daniel Lokshtanov, Dániel Marx, Paul Medvedev, Daniel Meister, George B. Mertzios, Martin Milanic, Neeldhara Misra, Haiko Müller, André Nichterlein, Rolf Niedermeier, Charis Papadopoulos, Christophe Paul, Daniel Paulusma, Marcin Pilipczuk, Michał Pilipczuk, Andrzej Proskurowski, Dieter B. Rautenbach, Saket Saurabh, Dimitrios M. Thilikos, Erik Jan van Leeuwen, Kristina Vušković

A surprisingly high number of the interesting computational problems arising from theory and applications can be formulated as *graph modification* problems. Here we are given as input a graph G, and the goal is to apply certain operations on G (such as vertex deletions, edge deletions, additions or contractions) in order to obtain a graph H with some particular property. For an example the classical Vertex Cover problem can be formulated as trying to change G into an edgeless graph by performing the minimum possible number of vertex deletions. The Cluster Editing problem is to change G into a disjoint union of cliques with a minimum number of edge deletions or additions. Graph modification problems have been studied quite extensively, and both algorithms for these problems and structural aspects have been thoroughly explored.

Graph modification problems have received a significant amount of attention from the perspective of Parameterized Complexity. In parameterized complexity input comes with a parameter k and the goal is to design *fixed parameter tractable* algorithms, i.e. algorithms with running time $f(k)n^{O(1)}$ for some, hopefully not too fast growing function f. The parameter k can be the size of the solution sought for, or it could be a number describing how structured the input instance is. For an example k could be the *treewidth* of the input graph. Over the last few years, our understanding of the parameterized complexity of graph modification problems has greatly improved. Fixed parameter tractable algorithms have been found for a number of fundamental graph modification problems. For several problems, surprising new algorithms with *subexponential* ($2^{o(k)}$) dependence on k have been developed.

There is a strong connection between graph modification problems and *graph classes*. A graph class is simply a set of graphs satisfying some common properties. Thus many, if not all, graph modification problems can be phrased as modifying the input graph G by as few operations as possible to make it fit into a particular graph class. There is a large and active Graph Classes research community that primarily investigates how restricting the *input* graph to a particular graph class affects the computational complexity of computational problems. In the setting of graph modification problems we have no restrictions on the input graph, but the problem definitions dictate which graph class the *output* graph should belong to. The main objective of the seminar was to bring together experts within Parameterized Algorithms and experts within Graph Classes to join forces on graph modification problems. We also invited experts from related areas, such as Structural Graph Theory and Bioinformatics. Structural graph theory, in order to learn of the new powerful graph theoretic tools being developed, and hopefully to apply them on graph modification problems. Bioinformatics, in order to better understand the relationship between the idealized models we study and real-world applications of graph algorithms.

The scientific program of the seminar consisted of 21 talks. 4 of these talks were longer (45 or 90 minute) presentations covering some of the most exciting developments on graph modification problems and related areas. We had one long talk for each of the main topics covered by the seminar. On Monday, Marcin and Michał Pilipczuk gave a joint 90 minute talk ("Subexponential parameterized complexity of completion problems") on parameterized algorithms. On Tuesday, Paul Medvedev gave a 45 minute talk ("An introduction to genome assembly and its relation to problems on graphs") showcasing how graph algorithms can be used in Bioinformatics applications. On Wednsday, Kristina Vušković gave a 45 minute presentation ("Weighted Independent Set in bull-free graphs") about how deep structure theorems can be useful in algorithm design, and on Thursday, Andreas Brandstädt gave a presentation ("Clique separator decomposition for a subclass of hole-free graphs") on graph classes. We believe that the invited talks were a good starting point for cross-community collaboration. The remaining

talks were 30 or 35 minute presentations on recent research of the participants. We made a point out of having fewer short talks, in order to leave more time for individual discussions and collaboration in groups, as well as for open problem sessions. The idea was to reserve almost all of the time between lunch and dinner for research. This was very well received by the participants. There were 3 fruitful open problem sessions, on Monday, Tuesday and Thursday. Notes on the presented problems can be found in this report.

4.12 New Perspectives in Shape Analysis

Organizers: Michael Breuß, Alfred M. Bruckstein, Petros Maragos, and Stefanie Wuhrer

Seminar No. 14072

Date: February 9–14, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.60

Participants: Mirela Ben-Chen, Sébastien Bougleux, Michael Breuß, Alex M. Bronstein, Michael M. Bronstein, Alfred M. Bruckstein, Leila De Floriani, Anastasia Dubrovina, Maurizio Falcone, Anna Hilsmann, Yong Chul Ju, Margret Keuper, Ron Kimmel, Iasonas Kokkinos, Aritom Kovnatsky, Petros Maragos, Gerlind Plonka-Hoch, Martin Rumpf, Frank R. Schmidt, Carola-Bibiane Schönlieb, William Smith, Sibel Tari, Silvia Tozza, Matthias Vestner, Martin Welk, Stefanie Wuhrer, Hao Zhang, Jovisa Zunic

Dagstuhl seminar 14072 *New Perspectives in Shape Analysis* took place February 9–14, 2014. 28 researchers from North America and Europe discussed state-of-the-art, current challenges, and promising future research directions in the areas of 2-D and 3-D shape analysis from a cross-disciplinary point of view. Participants included international experts from the fields of continuous-scale shape analysis, discrete shape analysis and sparsity, and numerical computing. The seminar consisted of an opening session, 11 scientific presentation sessions, as well as a break-out session, which provided room for in-depth discussions in small groups. Furthermore, there was time for extensive discussions both between the talks and in the evenings.

This seminar was motivated by the observation that in everyday life, geometric shapes surround us, and that the understanding of concepts describing these shapes is at the heart of various applications, such as ergonomic design, virtual shopping, scientific and medical visualization, realistic simulation, the design of natural user interfaces, and semantic scene understanding. Traditionally, the notion of shape has been studied either by analysing a sparse set of marker positions on 3-D shapes, mostly for medical imaging applications, or by analysing projections of shapes in 2-D images, mostly for image processing and computer vision applications. New challenges in the analysis and processing of such data arise with the increasing amount of data captured by sensors used to acquire shapes, and with modern applications such as natural user interfaces that require real-time processing of the input shapes. Recently, it has become increasingly affordable to digitise 3-D shapes using multiple modalities, such as laser-range scanners, image-based reconstruction systems, or depth cameras like the Kinect sensor. Using these dense 3-D shapes in the above mentioned applications requires processing and describing the shapes in an efficient and informative way.

The purpose of this seminar was to address these challenges with the latest tools related to geometric, algorithmic and numerical concepts. To do so, we brought together researchers working on shape analysis topics from different perspectives.

As the analysis of 3-D shapes and deformable shape models have received much interest recently, classic shape analysis tools from differential geometry have a fresh influence in the field. Being related to the issue how to represent shapes efficiently, the research areas of sparse data representation and machine learning have begun to influence shape analysis modelling and the numerics. Especially in the context of three-dimensional data (or even higher-dimensional data sets), efficient optimization methods will certainly become increasingly important since many shape analysis models can be cast in the form of an optimization problem.

While the fields of modelling and numerical computing are strongly related when it comes to shape analysis applications, modelling is seen as a hot topic in computer science while numerical computing is often seen as a mathematical domain. The purpose in bringing together researchers from those different communities sharing substantial interest in shape analysis was to explore the benefits of a *cross-disciplinary* point of view. More specifically,

- researchers in continuous-scale shape analysis brought to the meeting their knowledge of differential and variational models and also of classic numerical methods in the field,
- researchers in discrete shape analysis and sparsity brought to the meeting their knowledge about the latest techniques in efficient data representations and related machine learning techniques, as well as efficient data structures and discrete optimization methods, and
- researchers in numerical computing brought to the meeting their knowledge of numerical techniques for PDEs and optimization.

As the demands in the individual fields are high, the research groups in which the most interesting techniques are proposed are

quite specialised. This not only holds for discrete and continuous-scale modelling and numerical computing, but also for the areas of sparsity and machine learning that were discussed during this seminar. Because of this, there is no regular conference or workshop that serves as a meeting place for an exchange of ideas of these groups.

Promising new ways to combine the latest techniques from these different fields were identified during in-depth discussions in small groups. Some especially promising research directions in the areas of intrinsic structure detection, co-segmentation of shapes, shape from shading, modelling deformable shapes, and models for face shapes, were discussed in small groups during the break-out session.

4.13 Robots Learning from Experiences

Organizers: Anthony G. Cohn, Bernd Neumann, Alessandro Saffiotti, and Markus Vincze

Seminar No. 14081

Date: February 17–21, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.79

Participants: Michael Beetz, Sven Behnke, Alexandre Bernardino, Mustafa Blerim, Richard Bowden, Ivan Bratko, Francois Bremond, Anthony G. Cohn, Luc De Raedt, Krishna Sandeep Reddy Dubba, Martin Günther, Manfred Hild, Vaclav Hlavac, David C. Hogg, Lothar Hotz, Lorenzo Jamone, Stefan Konecny, Marek S. Kopicki, Jos Lehmann, Ales Leonardis, Masoumeh Mansouri, Ralf Moeller, Bernd Neumann, Davide Nitti, Laurent Orseau, Pierre-Yves Oudeyer, Federico Pecora, Sebastian Rockel, Alessandro Saffiotti, Luis Seabra Lopes, Muralikrishna Sridhar, Luc Steels, Sebastian Stock, Georgi Stojanov, Aryana Tavanai, Carme Torras, Emre Ugur, Markus Vincze, Jure Zabkar, Jianwei Zhang, Michael Zillich

Topics and Motivation

The ability to exploit experiences is an important asset of intelligent beings. Experiences provide a rich resource for learning, solving problems, avoiding difficulties, predicting the effects of activities, and obtaining commonsense insights. Current robots do not in general possess this ability, and this is a decisive reason for the often perceived "lack of intelligence" of current robotic systems: they repeat mistakes, do not learn to anticipate happenings in their environment, and need detailed instructions for each specific task.

Consider an everyday task of a service robot, such as grasping a cup from a cupboard and bringing it to a person sitting at a table. This task may occur in many variations and under unpredictable circumstances. For example, persons may sit at different sides of a table, a direct path to the table may be blocked, the table may be cluttered with various objects, hot water may be ready or not, the cup on the shelf may be upside-down, etc. It is clearly infeasible to provide the robot with precise instructions for all contingencies at design time or to specify tasks with highly detailed instructions for each particular concrete situation which may arise. Hence without such knowledge, robot behaviour is bound to lack robustness if the robot cannot autonomously adapt to new situations.

How would the robot, for example, avoid pouring coffee into an upside-down cup? Based on experiences with multiple pouring actions, the robot will have formed a conceptualisation of all concomitant circumstances of successful pouring, for example to pour into a "container". The robot may not know the name of this conceptualisation but will know that it must be open on top, hollow, empty, etc. Similarly, the robot may have encountered upside-down objects before and hence be able to conceptualise the corrective action of turning an object to make it a usable container.

This seminar has brought together experts and scholars from the robotics, learning, and knowledge representation communities to discuss current approaches to make robots learn from experiences. Emphasis was on the representation of real-world experiences and on exploiting experiences for autonomous acting in a changing or partially unknown environment.

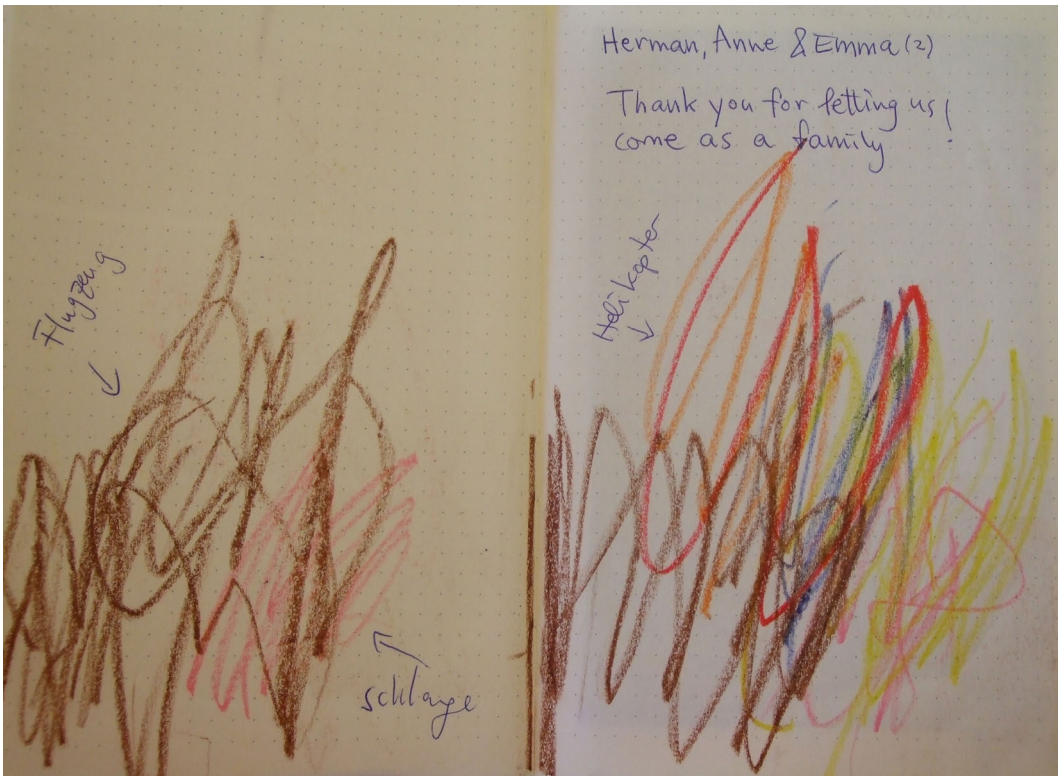

Fig. 4.5
Drawing for the Dagstuhl children's guest book by Emma (2), daughter of Dagstuhl Seminar 14091 participants Anne Dreimel and Herman Haverkort.

4.14 Visualization and Processing of Higher Order Descriptors for Multi-Valued Data

Organizers: Bernhard Burgeth, Ingrid Hotz, Anna Vilanova Bartroli, and Carl-Fredrik Westin

Seminar No. 14082

Date: February 16–21, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.110

Participants: Bernhard Burgeth, Guoning Chen, Tom Dela Haije, Çağatay Demiralp, Luc Florack, Andrea Fuster, Christoph Garth, Hans Hagen, Ingrid Hotz, Hans Knutsson, Max Langbein, Georgeta Elisabeta Marai, Rodrigo Moreno, Lauren O'Donnell, Evren Özarslan, Renato Pajarola, Ofer Pasternak, Jos B.T.M. Roerdink, Gerik Scheuermann, Marc Schöneich, Thomas Schultz, Nir Sochen, Markus Stommel, Anna Vilanova, Carl-Fredrik Westin, Eugene Zhang

Higher Order Descriptors for Multi-Valued Data

This seminar is the 5th in a series of Dagstuhl Seminars devoted to the visualization and processing of higher-order descriptors, of which tensors are a special case. They provide a natural language to describe phenomena in physics or image processing, e.g. medical imaging, fluid dynamics, or structural mechanics. Due to the increasing complexity of data generated in the engineering industry and the rapid advances in medical imaging, multi-valued data have gained significant importance in scientific visualization and image processing. Compared to their importance analysis and processing tools are still relatively rare.

In many scientific and engineering applications, as modern product development processes, simulations are an essential part of the advancement of the field. The results are large and complex data sets often comprising multi-filed data of various kind. Thereby, the tensor concept is essential for the description of physical phenomena related to anisotropic behavior. Examples for second-order tensors are stress and strain fields, inertia tensors, and orientation distribution tensors. Higher-order tensors occur when multiple vector fields or second-order tensors are set into relation, e.g., the stiffness tensor relating stress and strain. The field of engineering faces many open problems in tensor field analysis and visualization.

In medical imaging, multi-valued data include diffusion-weighted magnetic resonance imaging (dMRI), a medical imaging modality that allows the measurement of water diffusion in tissue (e.g., white matter or muscle) in vivo. Simple models approximate the diffusion in fibrous tissue by a second-order diffusion tensor (i.e., a positive semi-definite 3×3 matrix). But, often the acquired data is more complex and cannot be sufficiently described by the second-order tensor model and requires higher-order descriptors (i.e., higher-order tensors or spherical harmonics).

Even though these applications are very different in their nature they face many shared challenges associated with this highly complex data, which can profit from a multidisciplinary approach. The idea of this Dagstuhl Seminar was to bring together key researchers for disciplines ranging from visualization and image processing to applications in structural mechanics, fluid dynamics, and numerical mathematics.

Seminar Topics and Breakout Sessions

The seminar has been organized in presentation and breakout sessions. The presentation sessions gave the participants the possibility to presented recent developments in the multidisciplinary field. The talks covered a broad variety of topics related to both theoretical and practical issues. They served as basis for inspiring discussions across the application areas, which demonstrated that there are many shared issues related to analyzing and visualizing fields of tensors and higher-order descriptors. Besides the presentations, we put an emphasis on breakout sessions, which were very successful already in the previous meeting. They were targeted at fostering focused discussions in smaller groups. During a first session the group defined some driving objectives that partially already emerged in preparatory discussions:

- Statistics on higher-order descriptors and visualization of uncertainty
- Generalization of mathematical framework to higher-order descriptors
- Features on tensor visualization
- Next generation diffusion MRI

Most issues identified in the call have been discussed. Subjects that found special attention can be summarized as: Fundamental general topics, as tensor interpolation, statistics, morphology, and topology; Questions related to pattern descrip-

tion and detection; More specific issues like the analysis of ensembles, the visualization and measurement of differences and anomalies for engineering as well as for medical data sets. Further, there has been much interest in double pulse field gradient methods that have been discussed as possible next generation diffusion MRI. The outcome of the sessions can also be seen at http://www.dagstuhl.de/wiki/index.php/14082#Breakout_topics

The breakout sessions again turned out to be very successful. The format of the breakout sessions fits very well in the Dagstuhl environment promoting discussions and interactions. We could also observe that, for some topics, it is not easy to go beyond a list of challenges in such a short time frame. This motivates to strengthen these sessions by pre-defining topics in preparation of the meeting and asking selected participants for related statements.

Outcomes

The participants all agreed that the meeting was successful and stimulating. Seminar participants are already collaborating on a Springer book summarizing the results of the meeting. The Springer book will have about seventeen chapters authored by the meeting participants. It is also planned to summarize the results of two breakout sessions as a chapter of the book. The participants Thomas Schultz and Ingrid Hotz agreed on taking the lead for the collection of the contributions and the assembly of the book. We expect the book to be published in 2015.

It was voted that the group will apply for another meeting in this series, and that in addition to the current organizer (Ingrid Hotz) there will be two new organizers (Thomas Schultz, University Bonn Germany and Evren Özarslan, Harvard Medical School – Boston, US) for the next event.

Acknowledgement

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

4.15 Data Structures and Advanced Models of Computation on Big Data

Organizers: Alejandro López-Ortiz, Ulrich Carsten Meyer, and Robert Sedgewick

Seminar No. 14091

Date: February 24–28, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.129

Participants: Peyman Afshani, Deepak Ajwani, Lars Arge, Martin Aumüller, Hannah Bast, Timo Bingmann, Gerth Stølting Brodal, Andrej Brodnik, Martin Dietzfelbinger, Anne Driemel, Rolf Fagerberg, Johannes Fischer, Mordecai Golin, Torben Hagerup, Tanja Hartmann, Herman J. Haverkort, Meng He, John Iacono, Riko Jacob, Tsvi Kopelowitz, Moshe Lewenstein, Alejandro Lopez-Ortiz, Kurt Mehlhorn, Ulrich Carsten Meyer, Gabriel Moruz, Ian Munro, Markus E. Nebel, Patrick K. Nicholson, Jesper A. Sindahl Nielsen, John Owens, Rasmus Pagh, Rajeev Raman, Alejandro Salinger, Peter Sanders, Anita Schöbel, Nicole Schweikardt, Robert Sedgewick, He Sun, Sharma V. Thankachan, Laura I. Toma, Jan Vahrenhold, Sebastian Wild, Katharina A. Zweig

A persistent theme in the presentations in this Dagstuhl seminar is the need to refine our models of computation to adapt to modern architectures, if we are to develop a scientific basis for inventing efficient algorithms to solve real-world problems. For example, Mehlhorn's presentation on the cost of memory translation, Iacono's reexamination of the cache-oblivious model, and Sanders' description of communication efficiency all left many participants questioning basic assumptions they have carried for many years and are certain to stimulate new research in the future.

Better understanding of the properties of modern processors certainly can be fruitful. For example, several presentations, such as the papers by Aumüller, López-Ortiz, and Wild on Quicksort and the paper by Bingmann on string sorting, described faster versions of classic algorithms that are based on careful examination of modern processor design.

Overall, many presentations described experience with data from actual applications. For example, the presentations by Driemel and Varenhold on trajectory data described a relatively new big-data application that underscores the importance and breadth of application of classic techniques in computational geometry and data structure design.

Other presentations which discussed large data sets on modern architectures were the lower bound on parallel external list ranking by Jacob, which also applies on the MapReduce and BSP models commonly used in large distributed platforms; and by Hagerup who considered the standard problem of performing a depth first search (DFS) on a graph, a task that is trivial in small graphs but extremely complex on "big data" sets such as the Facebook graph. He proposed a space efficient algorithm that reduces the space required by DFS by a $\log n$ factor or an order of magnitude on practical data sets.

Schweikardt gave a model for MapReduce computations, a very common computing platform for very large server farms. Salinger considered the opposite end of the spectrum namely how to simplify the programming task as to take optimal advantage of a single server which also has its own degree of parallelism from multiple cores, GPUs and other parallel facilities.

In terms of geometric data structures for large data sets Afshani presented sublinear algorithms for the I/O model which generalize earlier work on sublinear algorithms. Sublinear algorithms are of key importance on very large data sets, which are thus presumably unable to fit in main memory. Yet most of the previously proposed algorithms assumed that such large data sets were hosted in main memory. Toma gave an external memory representation of the popular quad tree data structure commonly used in computer graphics as well as other spatial data applications.

4.16 Digital Evidence and Forensic Readiness

Organizers: Glenn S. Dardick, Barbara Endicott-Popovsky, Pavel Gladyshev, Thomas Kemmerich, and Carsten Rudolph

Seminar No. 14092

Date: Februrary 23–28, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.2.150

Participants: Aaron Alva, Carsten Bormann, Joseph Cannatac, Raymond Choo, Glenn S. Dardick, Günther Diederich, Jos Dumortier, Barbara Endicott-Popovsky, Katrin Y. Franke, Felix C. Freiling, Stefanie Gerdes, Pavel Gladyshev, Babak Habibnia, Nils-Peter Hercher, Florian Junge, Thomas Kemmerich, Nicolai Kuntze, David Manz, Christian Moch, Carsten Momsen, Heiko Patzlaff, Corinne Rogers, Carsten Rudolph, Viola Schmid, Isabel Taylor, Lee Tobin, Hein Venter, Rhythm Suren Wadhwa, Nigel Wilson, Stephen Wolthusen

This summary briefly recapitulates the outcomes of our seminar on digital evidence and forensic readiness. The main focus of the seminar was to work on a common cross-discipline understanding of notions of digital evidence and forensic readiness. In particular, technical notions in the view of IT security experts and the legal view were considered. Furthermore, relevance of differences in jurisdictions in different countries was also discussed.

The participants of the seminar came from 4 continents (Europe, U. S., Africa and Australia) and 12 countries. The group was a mix of experts from digital forensics, IT security, cyber security, archival sciences, criminal law, civil law, and cyber law. Thus, all relevant disciplines for digital evidence and forensic readiness were represented in the seminar, creating a perfect group for the task, but also a challenging communication environment that required good leadership in the interaction and discussions.

The main focus of the seminar was to develop a common view on what exactly can be considered secure and admissible digital evidence. The seminar was a first attempt to achieve progress towards this goal and therefore, a comprehensive coverage of the topic was not to be expected. Nevertheless, the international interest in the topic as well as the intensive discussions in the seminar show the relevance of the topic. The results of the seminar identify open issues in the area of digital forensics, but also proposes first substantial steps in the direction of establishing strong and internationally useful notions for digital evidence and forensic readiness.

Initial talks and discussions quickly revealed some of the majour challenges:

- The growing variety of types of potential digital evidence increases the problem to define clear technical guidelines for the collection and evaluation of data records for forensic use. Examples include mobile devices, data stored and processed via cloud service, huge infrastructures with distributed data, or big data with many possible interpretations of data found.
- In many cases, digital evidence cannot be directly related to data on one device. In particular in cloud environments, stored data is distributed over different countries and digital processes easily cross borders. Thus, digital evidence becomes a cross-jurisdictional issue that needs rules on how to deal with differences and contradictions in jurisdiction.
- Teaching and education is another challenge. One cannot expect all lawyers, attorneys, or judges to become experts on technical issues. however, a basic understanding of the area of digital evidence is essential to be able to decide if expert witnesses are required and also to be able to achieve correct interpretations of the report by expert witnesses.
- forensic readiness can guide the development of systems that collect, store, and provide secure digital evidence. However, the applicability of forensically ready technical solutions is restricted by privacy and also economy. Here, processes need to be defined and adequate procedures and regulations (also internationally) need to be found.

Four discussion groups were formed in the seminar to discuss *digital forensic readiness processes and procedures for investigators*, *notions of digital evidence*, *a forensic readiness landscape*, and *forensic readiness: evidence in a digital world*. More details of the results of the discussions in the working groups can be found in the sections below.

As one of the major results of the seminar can be identified that all participants understood and agreed on the need to initiate future research activities in the area of digital evidence and forensic readiness. The results also clearly show that this research must be international and inter-disciplinary. Furthermore, the seminar has proven that technically oriented IT security experts and experts from law can co-operate to advance the state of the art. The seminar has established new inter-disciplinary and international contacts that are suitable to build a new community that will drive this strand of work in the field of forensic readiness.

4.17 Preference Learning

Organizers: Johannes Fürnkranz, yke. Hüllermeier, Cynthia Rudin, Scott Sanner, and Roman Słowiński

Seminar No. 14101

Date: March 2–7, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.1

Participants: Nir Ailon, Fabio Aiolli, Antti Airola, Cédric Archambeau, Daniel Baier, Jerzy Błaszczyński, Róbert Busa-Fekete, Weiwei Cheng, Yann Chevaleyre, Krzysztof Dembczyński, Sébastien Destercke, Ad J. Feelders, Johannes Fürnkranz, Andreas Geyer-Schulz, Joachim Giesen, Salvatore Greco, Willem J. Heiser, Eyke Hüllermeier, Dietmar Jannach, Ulrich Junker, Kristian Kersting, Wojciech Kotłowski, Jérôme Lang, Eneldo Loza Mencá, Jérôme Mengin, Vincent Mousseau, Ingrid Oliveira de Nunes, Alena Otto, Tapio Pahikkala, Marc Pirlot, Michael Rademaker, Francesca Rossi, Scott Sanner, Michele Sebag, Eric Sibony, Roman Słowiński, Alexis Tsoukiàs, Nicolas Usunier, K. Brent Venable, Paolo Viappiani, Peter Vojtáš, Toby Walsh, Paul Weng, Christian Wirth

The topic of "preferences" has recently attracted considerable attention in Artificial Intelligence (AI) research, notably in fields such as autonomous agents, non-monotonic reasoning, constraint satisfaction, planning, and qualitative decision theory. Preferences provide a means for specifying desires in a declarative way, which is a point of critical importance for AI. Drawing on past research on knowledge representation and reasoning, AI offers qualitative and symbolic methods for treating preferences that can reasonably complement hitherto existing approaches from other fields, such as decision theory. Needless to say, however, the acquisition of preference information is not always an easy task. Therefore, not only are modeling languages and suitable representation formalisms needed, but also methods for the automatic learning, discovery, modeling, and adaptation of preferences.

It is hence hardly surprising that methods for learning and constructing preference models from explicit or implicit preference information and feedback are among the very recent research trends in disciplines such as machine learning, knowledge discovery, information retrieval, statistics, social choice theory, multiple criteria decision making, decision under risk and uncertainty, operations research, and others. In all these areas, considerable progress has been made on the representation and the automated learning of preference models. The goal of this Dagstuhl Seminar was to bring together international researchers in these areas, thereby stimulating the interaction between these fields with the goal of advancing the state-of-the-art in preference learning. Topics of interest to the seminar include

- quantitative and qualitative approaches to modeling preference information;
- preference extraction, mining, and elicitation;
- methodological foundations of preference learning (learning to rank, ordered classification, active learning, learning monotone models, ...)
- inference and reasoning about preferences;
- mathematical methods for ranking;
- applications of preference learning (web search, information retrieval, electronic commerce, games, personalization, recommender systems, ...).

The main goal of the seminar was to advance the state-of-the-art in preference learning from a theoretical, methodological as well as application-oriented point of view. Apart from that, however, we also hope that the seminar helped to further consolidate this research field, which is still in an early stage of its development. Last but not least, our goal was to connect preference learning with closely related fields and research communities (cf. Figure 4.6).

In order to achieve these goals, the program featured the following components:

- Monday was filled with 6 tutorial-type introductory talks about the use of preferences and the view on preference learning in the areas of machine learning, recommender systems, multi-criteria decision making, business and economics, artificial intelligence, and social choice, with the goal of familiarizing the members of the different communities with the basics of the other fields.
- Ten sessions were devoted to contributed presentations, each one with enough extra time for discussion. In case we ran over time, we gave priority to discussions. We were also able to flexibly integrate a few impromptu talks by participants.
- Two discussion sessions on Tuesday and Thursday afternoon were devoted to discussion how to establish closer connections between the different research areas that participated in this seminar.
- Wednesday afternoon featured a hike and an excursion to Trier with some wine tasting.

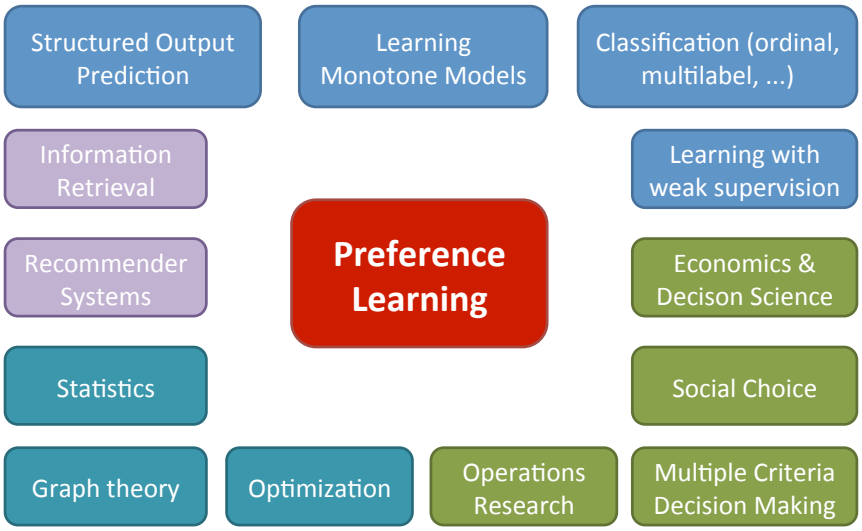

Fig. 4.6
Preference learning and related research areas within machine learning (blue), information retrieval (purple), applied mathematics (turquoise), and the decision sciences (green).

4.18 Combinatorics and Algorithmics of Strings

Organizers: Maxime Crochemore, James Currie, Gregory Kucherov, and Dirk Nowotka

Seminar No. 14111

Date: March 9–14, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.28

Participants: Dany Breslauer, Julien Cassaigne, Julien Clément, Maxime Crochemore, James D. Currie, Volker Diekert, Gabriele Fici, Johannes Fischer, Dominik D. Freydenberger, Anna E. Frid, Paweł Gawrychowski, Amy Glen, Štěpán Holub, Artur Jeż, Juha Kärkkäinen, Juhani Karhumäki, Steffen Kopecki, Gregory Kucherov, Manfred Kufleitner, Gad M. Landau, Alessio Langiu, Thierry Lecroq, Moshe Lewenstein, Florin Manea, Giancarlo Mauri, Robert Mercaş, Fillippo Mignosi, Mike Müller, Dirk Nowotka, Wojciech Plandowski, Ely Porat, Svetlana Puzynina, Antonio Restivo, Eric Rowland, Wojciech Rytter, Aleksi Saarela, Arseny M. Shur, Jamie Simpson, German Tischler, Esko Ukkonen, Mikhail V. Volkov

Processing strings efficiently is of concern in practically every application field. Understanding the combinatorial properties of sequences is a prerequisite for designing efficient algorithms on them. The Dagstuhl seminar 14111 has been concerned with exactly that: *Combinatorics and Algorithmics of Strings*.

This Dagstuhl seminar was attended by 41 researchers from 12 countries representing the two fields, algorithmics and combinatorics, about equally, although it needs to be mentioned that the overlap of these two communities is rather large. Inviting these close communities to Dagstuhl gave us the opportunity to start from substantial common ground and to work on scientific problems right from the beginning. Given that background, tutorials or other introductory sessions were not considered to be suitable elements for this seminar. Instead, much time was spent for problem posing and solving sessions. This seminar has clearly been research oriented.

The first seminar day, Monday, was entirely devoted to posing open problems. Based on those, the participants were able to form interest groups and engage into research activities early on. In the next days regular research talks and some more open problems were presented. However, time slots for research work were also allocated. On the last day of the seminar, Friday, we were able to already present some solutions to the problems posed in the beginning. In general, it is not to be expected that research problems are solved within a week (and most weren't), but it illustrates the impact of the meeting on catalysing research and collaboration between the participants.

The following two are great examples of such collaboration. Florin Manea asked about the complexity of deciding whether or not two words u and w are k-binomial equivalent, that is, is the number of occurrences of all scattered subwords up to length k equal in u and w? Contributions by Paweł Gawrychowski (polynomial Monte-Carlo algorithm in the logarithmic word-size RAM model), Juhani Karhumäki, and Wojciech Rytter (polynomial time on a unit-cost RAM model), and discussions with Dominik Freydenberger and Manfred Kufleitner finally led to the conclusion that the problem can be solved in polynomial time in the logarithmic word-size RAM model. Another problem was posed by Juhani Karhumäki and Michaël Rao (not present at the seminar) on the avoidability of shuffle squares. They asked: Does there exist an infinite word over some finite alphabet which avoids all factors that are a shuffle product of a word with itself? James Currie realized that shuffle squares can indeed be avoided applying the Lovász Local Lemma in his argument. However, this solution of avoidability in principle led to a proof for a very large alphabet, the size of which being a number of more than 40 digits. A few days after this Dagstuhl seminar Mike Müller improved that result by giving a rather low upper bound on the alphabet size of 10 on which shuffle squares can be avoided using a resent result by Joseph Miller. In general, it has to be noted that progress was made in many more areas and several papers in preparation were announced already.

Another notable highlight of the seminar was a session dedicated to word equations. Senior researchers of that particular research area, like Wojciech Plandowski and Volker Diekert, and young protagonists, like Aleksi Saarela, Štěpán Holub, and Artur Jeż, who talked about their recent efforts in developing the field, contributed and exchanged ideas. Such a unique assembly of major experts in word equations and their contributions at Dagstuhl was rather unique and a remarkable event.

In the light of such developments, it can be safely claimed that this seminar was a great success. Given the quality of presentations on this seminar and the constructive intensity of discussions, it is self-evident that a follow-up should be organised. We are grateful to all participants for their contributions to this successful seminar as well as to the staff of Schloss Dagstuhl for their great service.

4.19 Massive Open Online Courses: Current State and Perspectives

Organizers: Pierre Dillenbourg, Claude Kirchner, John C. Mitchell, and Martin Wirsing

Seminar No. 14112

Date: March 10–13, 2014 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.4.3.47

Participants: François Bry, Yiwei Cao, Carlos Delgado Kloos, Pierre Dillenbourg, Yannis Dimitriadis, Jens Dittrich, Gilles Dowek, Douglas H. Fisher, Antoine Flahault, Armando Fox, Serge Garlatti, David Glance, Jörg M. Haake, Jeff Haywood, Heinz Ulrich Hoppe, Claude Kirchner, Alain Mille, John C. Mitchell, Mike Sharples, Marcus Specht, Armin Weinberger, Martin Wirsing, Volker Zimmermann

Online education is not new; Massively Open Online Courses (MOOCs) are. Their uniquely powerful combination of classical digital teaching tools (videos, audios, graphics or slides), individualized tools for acquiring and validating knowledge, and appropriate use of dedicated social networks makes them a new and formidable means of accessing knowledge and education. If backed up with scientific and pedagogical excellence, MOOCs allow one to reach and teach simultaneously tens of thousands and even hundreds of thousand of learners in a new pedagogical dynamic.

Of the numerous MOOCs initiatives that have recently emerged, especially in the US and Europe, a few seem to be surfacing with an extremely important impact. This creates a very new situation and indeed can be considered as the informatics community's first main impact on knowledge dissemination and teaching. MOOCs will very likely induce a radical change in teaching mechanisms and their links to the economic and production systems. The consequences with respect to the transmission of culture and educational content, and on society as a whole, will be deep.

This situation raises many questions in a range of different disciplines with respect to ethics, intellectual properties, and data protection and privacy, necessitating an in-depth understanding of the current state of affairs and future trends in these research areas.

This Dagstuhl Perspectives Workshop brought together leading researchers and practitioners working in or on MOOCs initiatives in order to provide a forum for discussing participants' current experiences and initial feedback. Scientists from several key disciplines, including informatics, pedagogy, economy, psychology and sociology, have meet to discuss the current state of the situation and envision the next steps. In particular, they have addressed questions relative to current research on the pedagogical engineering of MOOCs, economical models, ethical issues, the technical development of platforms, and sharing.

The first day of the workshop consisted of a series of presentations in which each participant presented those topics and developments he or she considered most relevant for the future development of MOOCs. On the second and third day the participants divided into several working groups according to the main thematic areas that had been identified on the first day.

From the working groups outcomes, a Manifesto has been worked out and will appear in the *Dagstuhl Manifestos Collection*.[14]

[14] The manifesto was published in October 2014 and is available at http://dx.doi.org/10.4230/DagMan.4.1.1.

4.20 Computational Complexity of Discrete Problems

Organizers: Anna Gal, Michal Koucký, Oded Regev, and Rüdiger Reischuk

Seminar No. 14121

Date: March 16–21, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.62

Participants: Eric Allender, Eli Ben-Sasson, Beate Bollig, Mark Braverman, Harry Buhrman, Amit Chakrabarti, Arkadev Chattopadhyay, Irit Dinur, Lance Fortnow, Anna Gál, Oded Goldreich, Kristoffer Arnsfelt Hansen, Prahladh Harsha, Johan Hastad, Valentine Kabanets, Subhash Khot, Hartmut Klauck, Swastik Kopparty, Michal Koucký, Matthias Krause, Meena Mahajan, Pierre McKenzie, Peter Bro Miltersen, Ankur Moitra, Pavel Pudlák, Jaikumar Radhakrishnan, Anup Rao, Oded Regev, Rüdiger Reischuk, Michael Saks, Rahul Santhanam, Shubhangi Saraf, Nicole Schweikardt, Ronen Shaltiel, Amir Shpilka, Amnon Ta-Shma, Till Tantau, Thomas Thierauf, Christopher Umans, Dieter van Melkebeek, Omri Weinstein, David P. Woodruff, Amir Yehudayoff

Introduction and goals

Computational complexity aims to answer what is efficiently solvable and what is not on various computational models. This means providing upper and lower bounds on the necessary resources such as time, space or communication, and establishing connections among the different models.

There are intricate connections between complexity measures in different computational models. For instance, circuit size is closely related to computation time, whereas circuit depth and branching program size are closely related to computation space. Breaking the current barriers of our understanding in any of these models would have major consequences in several of the other models as well.

Investigating the connections between the various computational models and subareas of computational complexity has already led to many exciting results. In recent years several novel techniques have been introduced in computational complexity, resulting in a number of breakthroughs, some of which are still actively investigated. In particular, information-theoretic techniques have led to tremendous progress in our understanding of communication complexity, such as new methods to compress interactive communication, very efficient ways to immunize protocols against corruption of the communication by an adversary, and a better understanding of so-called direct product questions. This progress in turn led to progress in our understanding of the streaming model, in which one needs to process massive amounts of received data without being able to store it. Semi-definite programming, a technique originally used in optimization and in the design of approximation algorithms, has led to a tight and very elegant characterization of quantum query complexity. In the area of hardness of approximation, new approaches to prove the Unique Game Conjecture (which is one of the most central open questions in the area) have been suggested. Finally, a recent breakthrough separation of the class NEXP (nondeterministic exponential time) from the class ACC^0 (bounded depth circuits with counting) rests on a new technique that derives a lower bound for a non-uniform model from an upper bound on satisfiability in the uniform setting; this technique opens up a new range of possible connections between uniform and non-uniform models.

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

Organization of the meeting

43 researches from around the world participated in the seminar including a substantial number of young researchers. Each day, Monday to Thursday, we started by a longer talk surveying recent results in specific areas that were chosen beforehand. We had the following survey talks:

- Shubhangi Saraf: Recent developments in arithmetic circuits
- Subhash Khot: On the unique games conjecture and the approximation resistance of predicates.
- Mark Braverman: Recent progress on interactive error correction: an overview.
- Ankur Moitra: Extended formulations and information complexity.

Additionally, on Friday we started with a survey on recent progress in algorithms for matrix multiplication presented by Chris Umans. The tutorials were followed by shorter talks by

other participants. Afternoons were reserved for discussions in impromptu groups. In late afternoon on Monday, Tuesday and Thursday we had several additional short talks. On Wednesday evening we organized a rump session where everyone could present an open problem or announce a new result. One of the open problems from this session on the relationship between information cost and communication complexity presented by Omri Weinstein was very recently resolved.

Topics covered by the seminar

The talks of the workshop fit into several subareas of computational complexity. We summarize the talks next. Detailed abstracts of the talks can be found at the end of this report.

Circuit complexity

One of the goals in circuit complexity is to prove strong lower bounds on the size of circuits computing explicit functions. Even in the case of bounded depth circuits the known lower bounds deteriorate quickly with depth. Oded Goldreich discussed approaches to prove strong lower bounds of almost the type $2^{\Omega(n)}$ in such a setting by focusing on certain kinds of multilinear functions.

Another approach to proving lower bounds was presented by Anup Rao, who showed new lower bounds for bounded-depth circuits with arbitrary gates when the fan-in of gates is strictly smaller than n.

Valentine Kabanets considered the interplay between Boolean formulas and harmonic analysis of functions computed by Boolean formulas. He showed that functions represented by sub-quadratic formulas over the basis AND, OR and NOT have constrained Fourier coefficients. Among other things, this fact leads to new learning algorithms for such functions.

Shubhangi Saraf reviewed recent progress towards separating Valiant's classes VP and VNP, the arithmetic analogues of P and NP.

Eric Allender in his talk focused on another aspect of circuit complexity by providing improved upper bounds on the level of counting hierarchy in which certain problems involving arithmetic circuits lie.

Beside proving lower bounds several talks also focused on algorithmic aspects of circuits. Kristoffer Arnsfelt Hansen discussed the circuit complexity of several graph problems when the graphs have bounded cut-width, and Swastik Kopparty showed in his talk an efficient way of indexing irreducible polynomials over finite fields which may serve as a useful tool in designing efficient arithmetic circuits.

Amir Yehudayoff studied the growth rate of symmetric polynomials with possible applications in pseudorandomness.

Communication complexity and its applications

The classical theory of error correcting codes addresses mainly the question of one-way communication over unreliable channel. In communication complexity the main issue is to minimize the amount of communication between two interacting parties whose goal is to evaluate some joint function of their respective inputs. In this scenario the communication goes both ways. Mark Braverman gave a summary of results on error correcting techniques when the two parties communicate over unreliable channel.

Pavel Pudlák presented approaches to constructing good error correcting codes for interactive communication (so-called tree codes) based on properties of certain matrices.

Another popular research topic in communication complexity is information complexity. This topic was discussed by Omri Weinstein. He showed a new technique to estimate interactively the amount of information leaked by the two players about their inputs during a two party communication. This might have applications for secure communication.

Hartmut Klauck presented an interplay between quantum and classical communication, and established that in certain setting quantum communication can be replaced by classical messages.

Mike Saks provided a surprisingly simple protocol for certain class of functions in the number-on-the-forehead multi-party model.

Ankur Moitra presented a survey on recent results regarding extended formulation approach to solving hard combinatorial problems. In this context he also successfully applied techniques from communication complexity.

Communication complexity is a major tool in the analysis of data stream algorithms, algorithms that can process huge data sets while utilizing only little memory. David Woodruff presented a surprising fundamental result showing that a large class of streaming algorithms can be simulated using only linear sketches of the data stream. This could simplify design of data stream algorithms.

Amit Chakrabarti considered a model for processing large data streams with the help of an untrusted but powerful helper (e.g. cloud service). He discussed a relationship between this model and Arthur Merlin communication protocols.

Inapproximability

When we lack efficient algorithms for various problems that are NP-complete we may try to solve them approximately. In some cases, even that is hard as demonstrated by Prahladh Harsha in his talk on inapproximability of coloring of hypergraphs.

On the other hand, Johan Håstad presented a new algorithm for finding a satisfying solution to a CNF formula when all clauses in the formula can simultaneously be satisfied by majority of their literals. When the formula does not have such a property the problem becomes NP-complete.

Irit Dinur discussed her results on testing whether a given function is a direct product of some function with application to parallel repetition, and Eli Ben-Sasson explained his result on constructing linear-size probabilistically checkable proofs (PCP) that can be checked using n^ϵ queries.

Pseudorandomness

Construction of pseudorandom generators for Boolean circuits is currently reasonably well understood. However, in non-Boolean setting such as in the case of multi-output functions or arithmetic circuits we still lack good understanding of the problem. Ronen Shaltiel presented pseudorandom generators with optimal seed length for multi-output functions computed by polynomial size circuits, and Amnon Ta-Shma presented a new construction of hitting set generators for low-degree polynomials.

A central problem for which we know a very efficient randomized algorithm but no deterministic one is the problem of testing whether a polynomial is identically zero. Meena Mahajan considered the problem of testing whether a polynomial represented by an arithmetic formula that reads each variable at most three times is zero or not. She provided a deterministic algorithm for this problem.

Amir Shpilka presented a new algorithm for the closely related problem of testing whether two polynomials are the same up to a linear transformation of variables.

Other models

Harry Buhrman presented a new model of computation, catalytic space, in which in addition to the usual limited work space we have essentially unlimited amount of extra space which is however full of data that have to be preserved. He exhibited the surprising power of this extra space that allows one to compute functions that we do not know how to compute using only the limited work space.

Matthias Krause discussed the issue of cryptographic authentication by devices with limited resources which are not able to evaluate the standard cryptographic primitives like RSA and AES. He proposed solutions for those situations and reported on an actual implementation.

Jaikumar Radhakrishnan considered the bit-probe complexity of a data structure for storing sets (set-membership problem). He presented a very elegant and more efficient solution for this problem.

Thomas Thierauf presented an algorithm to compute the number of perfect matchings in K_5-free graphs. In general graphs this problem is considered to be hard.

Till Tantau talked about parallel algorithms in the context of fixed parameter tractability. He defined the notion and presented parallel algorithms in that context.

Chris Umans presented an overview of recent progress on matrix multiplication.

Conclusion

Understanding the computational complexity of various problems is the primary goal of theory of computing. In the past several years there has be tremendous progress in various areas of complexity for example, in communication complexity, arithmetic circuit complexity and derandomization. This progress brings us closer to the goal of understanding computation. Yet, as we have seen new relevant concepts and models emerge, e.g., information cost and catalytic computation. Despite all the progress that have been achieved since our previous meeting three years ago, and in the light of the new developments, there is a general consensus among the participants of the seminar that there is still long way ahead of us before we gain a good understanding of limits of efficient computation and resolve many of the central problems in computational complexity.

We like to thank the staff at Dagstuhl who – as usual – provided a marvelous surrounding to make this a successful meeting with ample space for undisturbed interactions between the participants.

Fig. 4.7
Jaques Gassmann – Diaspore. Part of the Dagstuhl art collection and donated by: Reinhard Wilhelm, Andreas Zeller, and the participants in Dagstuhl Seminars 07101, 08031, 08051, 07341, 08302, and 08441.

4.21 Verification of Cyber-Physical Systems

Organizers: Rupak Majumdar, Richard M. Murray, and Pavithra Prabhakar

Seminar No. 14122

Date: March 17–21, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.85

Participants: Erika Ábrahám, Calin A. Belta, Sergiy Bogomolov, Ken Butts, Xin Chen, Deepak D'Souza, Thao Dang, Jyotirmoy Deshmukh, Georgios Fainekos, Goran Frehse, Sebastian Gerwinn, Radu Grosu, Bruce H. Krogh, Mircea Lazar, Rupak Majumdar, Oded Maler, Ian M. Mitchell, Sayan Mitra, Richard M. Murray, André Platzer, Pavithra Prabhakar, Jean-François Raskin, Harald Ruess, Indranil Saha, Sriram Sankaranarayanan, Konstantin Selyunin, Danielle Tarraf, Ashish Tiwari, Mahesh Viswanathan

Introduction

Cyber-physical systems are systems in which there exists a tight coupling between computation, communication and control. The drastic reduction in the cost of sensing, actuating, computing and communicating technology has enabled the proliferation of this new genre of engineered systems in which a network of embedded processors interact tightly with the physical world to achieve complex functionalities. They have applications in a wide-range of systems spanning communication, infrastructure, energy, health-care, manufacturing, military, robotics and transportation.

Cyber-physical systems are believed to be the systems of the future with an impact on the engineering systems technology comparable to the impact the internet had on the information systems. Governments around the world have taken several initiatives to exploit this potential. The report of the US President's Council of Advisors on Science and Technology (PCAST) has placed Cyber-Physical Systems on the top of the priority list for federal research investment. The European Union has recognized the strategic importance of Embedded Computing Systems and has launched the ARTEMIS Joint Technology Initiative (JTI) as part of the FP7 program. Also, the latest European Commission Work Programme 2013 for Information and Communication technologies identifies this with the Objective ICT-2013.3.4 dedicated to Advanced Computing, Embedded and Control systems.

Cyber-Physical Systems have immense potential for a long-term impact on the society. At the same time, the unprecedented complexity arising due to the interleaving of the cyber and the physical components is overwhelming. On one hand, digital systems operate is a discrete manner, where computation and communication proceed in synchronization with the processor cycles. On the other hand, physical systems execute continuously in dense real-time. Hence, cyber-physical systems are complex systems exhibiting both discrete and continuous behaviors, and are networked and/or distributed with possibly humans in the loop. The grand challenge of the near future is the development of design methodologies and tools to cater to the development of reliable cyber-physical systems.

Model-based development has emerged as the de facto product development process in several domains including automotive and aeronautics. Here, the product development cycle begins with an abstract mathematical model of the system which is subject to rigorous analysis. The code is then generated from the model either automatically or manually. This enables early detection and correction of bugs which in turn results in the reduction of development costs and time, thereby providing companies with a competitive edge. However, the techniques used for analysis based on simulation of the mathematical models is still ad hoc, and does not provide the high level of reliability guarantees expected out of safety-critical CPS. Formal verification is an alternative approach which aims to provide a proof of correctness of the system. It is a promising technique for achieving the goal of developing high confidence cyber-physical systems.

Outcomes of the seminar

The seminar focused on the challenges in the application of formal methods towards verification of CPS. The seminar had a total of 28 participants with a mix of computer scientists and control theorists.

Tutorials

Given the cross disciplinary nature of the seminar, 6 tutorials were arranged on the following topics to provide a common ground to enable researchers with different backgrounds to communicate.

1. Simulation-Based Techniques for the Falsification of Cyber-Physical Systems

2. Verification of Automotive Engine Control
3. Formal Methods for Control Design
4. On Optimal and Reasonable Control in the Presencve of Adversaries
5. Compositionality Results for Cardiac Cell Dynamics
6. Logic of Hybrid Games

Sessions

The following topics were identified as important issues in the application of formal verification to CPS. A separate session was dedicated to discuss the topics in the context of CPS.

1. **Simulation based methods**: Application of simulation techniques for performing verification of CPS was discussed.
2. **Using verification for control design**: This session focused on the application of formal verification techniques such as those based on abstractions for control design.
3. **Foundation of CPS**: This session discussed the complexity and decidability of problems in verification and control of CPS.
4. **Applications**: This session discussed the methods and challenges in the verification of aircraft control, biological systems and multi-robot path planning.
5. **Abstractions**: This session discussed the issues regarding simplification techniques for scalable analysis of CPS.
6. **Lyapunov based methods**: This session discussed notions of stability and techniques for their analysis.
7. **Constraint solving**: Several verification problems can be formulated as constraint solving problems. This session discussed the challenges in constraint solving problems arising in CPS.
8. **Symbolic Verification**: This session discussed problems related to building efficient algorithms and tools for symbolic state-space exploration.

Research Directions

The seminar successfully fostered communication between computer scientist and control theorist. Some challenges and research directions were identified such as the need for the development of compositional reasoning of CPS with multiple components and lightweight analysis methods to boost scalability (such as using simulation for verification).

4.22 Computational Models of Cultural Behavior for Human-Agent Interaction

Organizers: Elisabeth André, Ruth Aylett, Gert Jan Hofstede, and Ana Paiva

Seminar No. 14131

Date: March 23–28, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.103

Participants: Jan M. Allbeck, Jens Allwood, Elisabeth André, Ruth Aylett, Nadia Berthouze, Timothy Bickmore, Emmanuel G. Blanchard, Justine Cassell, Nick Degens, Birgit Endraß, Colette Faucher, Samantha Finkelstein, Kobi Gal, Lynne Hall, Béatrice Hasler, Rüdiger Heimgärtner, Dirk Heylen, Gert Jan Hofstede, W. Lewis Johnson, Kristiina Jokinen, Catholijn M. Jonker, Rilla Khaled, Felix Kistler, Tomoko Koda, Sarit Kraus, Brigitte Krenn, H. Chad Lane, Lydia Lau, Jean-Claude Martin, Bilyana Martinovski, Samuel Mascarenhas, Phaedra Mohammed, Yukiko I. Nakano, Toyoaki Nishida, Hirotaka Osawa, Ana Paiva, Catherine Pelachaud, Matthias Rauterberg, Matthias Rehm, Kasper Rodil, Suleman Shahid, David R. Traum

The seminar was the first international meeting on computational models of culture. The gathering of a multi-disciplinary team of experts with a background in intelligent virtual agents, human-computer interaction and cultural theories provided us with a lot of inspirations for future research projects. We did not only identify relevant topics for a roadmap on computational models of culture, but also worked out a number of intriguing applications for cultural agents.

In order to exploit the momentum and enthusiasm generated by the seminar, we discussed some ideas for community building. We plan to organize a follow-up Dagstuhl seminar in about two or three years. However, some participants felt they would not like to wait for such a long time and spontaneously decided to have an additional workshop on Cultural Models for Intelligent Virtual Agents at the Fourteenth International Conference on Intelligent Virtual Agents (IVA 2014) in Boston this year. Furthermore, we discussed the edition of a book and/or a special issue. In addition, we talked about possibilities to share and distribute corpora to support comparative studies of culture-specific behaviors.

4.23 Interaction and Collective Movement Processing

Organizers: Maike Buchin, Luca Giuggioli, Marc van Kreveld, and Guy Theraulaz

Seminar No. 14132

Date: March 23–28, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.138

Participants: Ulf Blanke, Kevin Buchin, Maike Buchin, Edward A. Codling, Matt Duckham, Roland Geraerts, Luca Giuggioli, Asja Jelic, Andrea Kölzsch, Patrick Laube, Maarten Löffler, Ugo Lopez, Ran Nathan, Ko Nishino, Nicholas Ouellette, Andrea Perna, Ross Purves, Stef Sijben, Jack Snoeyink, Bettina Speckmann, Frank Staals, Greg Stephens, Guy Theraulaz, Gilles Tredan, Stefan van der Spek, Marc van Kreveld, Yusu Wang, Carola Wenk, Yossi Yovel

The Dagstuhl Seminar on Interaction and Collective Movement Processing brought together a group of 30 scientists with varied backgrounds, but with a shared interest in computations involved in the processing of moving entity data, like humans or animals. There are different reasons for such computations: they are needed for the initial processing (cleaning, recognition), for the analysis (derived properties, patterns), and for more advanced features like characterizing and modelling interaction between entities. This seminar focused on the latter, the hardest of these tasks. The majority of the participants had a background in ecology, behavioral sciences, or geometric algorithms, but there were also participants from statistical physics, GIScience, and computer vision.

The seminar featured four invited talks in four main research fields: ecology (Greg Stephens), computational geometry (Jack Snoeyink), GIScience (Patrick Laube), and collective motion (Andrea Perna). The remainder of the program consisted of short presentations, open problem sessions, break-out groups to work on open problems, and reporting sessions based on research done in the break-out groups.

While the original intention was to tackle the challenging problems of interaction and collective motion, part of the research was done on other closely related topics in movement analysis, like quality issues in movement analysis. The problems that were investigated – also described in this report – have led to the start of new research, which was exactly the purpose of the seminar.

The participants enjoyed both the seminar setting and the interdisciplinarity of the seminar, which gave a new impulse to the research of many. A number of collaborations have started up, and we hope that these not only lead to publications but also to longer lasting collaborations. While all participants would be happy to return to such a seminar later, it was agreed that the focus will shift to keep the dynamics and cross-fertilization of different research fields.

4.24 Reachability Problems for Infinite-State Systems

Organizers: Javier Esparza, Alain Finkel, Pierre McKenzie, and Joel Ouaknine

Seminar No. 14141

Date: March 30 to April 4, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.153

Participants: Mohamed Faouzi Atig, Parosh Aziz Abdulla, Amir M. Ben-Amram, Michael Blondin, Bernard Boigelot, Mikolaj Bojańczyk, Rémi Bonnet, Ahmed Bouajjani, Olivier Bournez, Tomas Brazdil, Véronique Bruyère, Michaël Cadilhac, Giorgio Delzanno, Laurent Doyen, Diego Figueira, Alain Finkel, Pierre Ganty, Stefan Göller, Jean Goubault-Larrecq, Christoph Haase, Peter Habermehl, Radu Iosif, Petr Jančar, Stefan Kiefer, Barbara König, Antonín Kučera, Slawomir Lasota, Ranko Lazić, Jérôme Leroux, Richard Mayr, Pierre McKenzie, Roland Meyer, Joel Ouaknine, Andreas Podelski, Igor Potapov, Fernando Rosa-Velardo, Sylvain Schmitz, Philippe Schnoebelen, Luc Segoufin, Jeffrey O. Shallit, Grégoire Sutre, Wolfgang Thomas, Igor Walukiewicz, James Worrell

Many standard verification problems can be rephrased as reachability problems, and there exist powerful methods for determining reachability in infinite-state systems. However, applications require not only decidability results, but provably optimal algorithms. The seminar focussed on complexity and algorithmic issues for the verification of infinite-state systems, with special emphasis on reachability problems.

Verification of finite-state systems can be illustrated by considering the case of counter systems, i. e., computational models combining a finite-state control with counters. Counter systems have been used to model distributed protocols, programs with recursive parallel threads, programs with pointers, broadcast protocols, replicated finite-state programs, asynchronous programs, etc. If zero-tests are allowed – one speaks of "Minsky machines" –, counter systems have the power of Turing machines, and so all their verification problems are undecidable. On the other hand, many problems can be decided when zero-tests are forbidden – one speaks of VASS, for "vector addition systems with states", or equivalently "Petri nets". In particular, reachability for VASS was shown decidable in 1982, and this can be leveraged into many more positive results. Moreover, researchers developed techniques that, while necessarily incomplete, allow analysing many questions: reversal-bounded analysis à la Ibarra, accelerations à la FAST, or well-structured extensions of VASS, see e. g., the forward analysis procedure. In turn, these techniques have led to many new theoretical results. For instance, it has been shown that the reachability sets of both reversal-bounded counter automata and flat counter automata are effectively definable in Presburger arithmetic (assuming some additional conditions).

The seminar addressed the following topics:

- Complexity of reachability on various models: parameterized counter systems, lossy channel systems, lossy counter systems, at counter systems, reversal-bounded counter systems, and other.
- Decidability and complexity of reachability problems for Petri nets extensions: timed Petri nets, Petri nets with one zero-test, with one unbounded counter, linear dynamical systems, BVASS, data nets, and other.
- Recent development and uses of the theory of well-structured transition systems.
- Decidability and complexity of reachability for systems with multiple (constraints) stacks: multiphase, reversal-bounded, and other.
- Games on infinite-state systems: counter automata, timed systems, weighted automata. Games with energy constraints.
- Monadic logics with costs.
- New developments in the algorithmics of Presburger logics; SMT-solvers.

4.25 Spatial Reference in the Semantic Web and in Robotics

Organizers: Aldo Gangemi, Verena V. Hafner, Werner Kuhn, Simon Scheider, and Luc Steels

Seminar No. 14142

Date: March 30 to April 4, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.3.181

Participants: Brandon Bennett, Mehul Bhatt, Saša Bodiroža, Martin V. Butz, Nicholas Chrisman, Helen Couclelis, Anne Cregan, Max J. Egenhofer, Andrew U. Frank, Christian Freksa, Aldo Gangemi, Giancarlo Guizzardi, Verena V. Hafner, Pascal Hitzler, Werner Kuhn, Bruno Lara Guzman, Vivien Mast, Mihoko Niitsuma, Kai-Florian Richter, Katharina J. Rohlfing, Yulia Sandamirskaya, Thomas Scharrenbach, Simon Scheider, Michael Spranger, Stephan Winter, Stefan Wölfl, Diedrich Wolter

Motivation

Places ("downtown"), spatial objects ("highway 1") and localized events ("hurricane Katrina"), are commonly referred to in the Semantic Web. They serve to search for and link to information across domains. Spatial reference systems, such as WGS84[15], allow for representing such references as points or regions. This makes them amenable not only for mapping, but also for powerful location-based querying, navigation support and computing.

Spatial references are also fundamental in embodied cognition and robotics. Egocentric and allocentric spatial reference frames underlie robot learning and interaction. Decades of research in cognitive robotics highlight the role of social interaction, joint attention, language games, and visual discrimination games in establishing referents for symbols. The most well-known experiment is that of the Talking Heads. Spatial relations, such as right, front, left, behind, serve to name and identify other objects in a self-organizing vocabulary. Affordance-based cognition is a source of spatial reference in robots as well as in humans. However, so far, this research is only loosely connected to information science and the Semantic Web.

Existing options to localize information in the Semantic Web and in Robotics through coordinate systems cover only limited cases of spatial reference. Humans localize referents in space in many ways, based on different tasks and spatial competencies. For example, the location of a workplace may be linked to people, tasks, and infrastructures. It can be specified in terms of a coordinate system or, alternatively, in terms of containment, connectedness and accessibility in a building; yet another option is to specify it by the possibility to perform certain activities, such as sitting or reading and writing at the workplace.<

The seminar

This Dagstuhl Seminar brought together leading international researchers from the Semantic Web, Spatial Cognition, Geo-informatics and Cognitive Robotics to work on the application, synthesis, formal construction, extension, and use of spatial reference systems, identifying challenges and research opportunities. The seminar gathered 27 researchers, 9 from Spatial Cognition and reasoning, 6 from Geo-informatics, 7 from Cognitive Robotics, and 5 from the Semantic Web.

Seminar participants identified a number of concrete links between these communities that are being exploited for future research and development. For example, spatial reference systems of robots and corresponding cognitive spatial concepts can be used in order to describe resources accessible in the world, and Semantic Web technology to publish those descriptions for information access. Locations can be described in ways which are more closely related to humans, based on qualitative relations and environmental referents, and for environments which are difficult to localize by a GPS. In this way, it becomes possible to share location descriptions among humans and robots and thus to localize resources of interest (e. g. rooms, people, places) published in the Web of data. Vice versa, spatial referents and descriptions in the Semantic Web may guide robots towards accessible things in the world. Robots may function as embodied surrogates of human observers exchanging information on the Web of Data encoded in terms of their own reference systems.

[15] The world geodetic system https://en.wikipedia.org/wiki/World_Geodetic_System, a standard reference ellipsoid and coordinate frame for encoding locations on the earth surface.

4.26 Evaluating Software Verification Systems: Benchmarks and Competitions

Organizers: Dirk Beyer, Marieke Huisman, Vladimir Klebanov, and Rosemary Monahan

Seminar No. 14171

Date: April 21–25, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.4.1

Participants: Aws Albarghouthi, Ezio Bartocci, Bernhard Beckert, Dirk Beyer, Marc Brockschmidt, David Cok, Gidon Ernst, Marie Farrell, Jean-Christophe Filliâtre, Bernd Fischer, Alberto Griggio, Radu Grigore, Arie Gurfinkel, Matthias Heizmann, Marieke Huisman, Bart Jacobs, Alexey Khoroshilov, Joseph Roland Kiniry, Vladimir Klebanov, K. Rustan M. Leino, Stefan Löwe, Antoine Miné, Rosemary Monahan, Wojciech Mostowski, Peter Müller, Petr Müller, Vadim Mutilin, Andrei Paskevich, Nadia Polikarpova, Arend Rensink, Philipp Rümmer, Andrey Rybalchenko, Gerhard Schellhorn, Markus Schordan, Carsten Sinz, Bernhard Steffen, Jan Strejcek, Michael Tautschnig, Mattias Ulbrich, Jaco van de Pol, Angela Wallenburg, Philipp Wendler, Thomas Wies

The seminar aimed to advance comparative empirical evaluation of software verification systems by bringing together current and future competition organizers and participants, benchmark maintainers, as well as practitioners and researchers interested in the topic.

The objectives of the seminar were to (1) advance the technical state of comparative empirical evaluation of verification tools, (2) achieve cross-fertilization between verification communities on common/related issues such as selection of relevant problems, notions of correctness, questions of programming language semantics, etc., (3) explore common evaluation of different kinds of verification tools, its appropriate scope and techniques, (4) raise mutual awareness between verification communities concerning terminology, techniques and trends, and (5) promote comparative empirical evaluation in the larger formal methods community.

Altogether, 43 researchers and practitioners have attended the seminar. A vast majority of the attendees (almost 90%) have participated either in the SV-COMP or in the VerifyThis (and related, e. g., VSComp/VSTTE) verification competitions. For lack of better terms, we tend to refer to these communities as the "automatic verification" and "deductive verification" community respectively, though, as these labels are based on pragmatics and history rather than on technical aspects.

The presentations, hands-on sessions, and discussions provided valuable feedback that will help competition organizers improve future installments. To continue the effort, a task force will compile a map of the – in the meantime very diverse – competition landscape to identify and promote useful evaluation techniques. It was agreed that evaluation involving both automatic and deductive verification tools would be beneficial to both communities as it would demonstrate the strengths and weaknesses of each approach. Both SV-COMP and VerifyThis will be associated with the ETAPS conference in 2015.

A call to the public: It has been reported that competition-verification challenges have been used as homework for students. The seminar organisers would appreciate feedback and experience reports from such exercises.

Seminar Structure

The seminar was structured as a highly interactive event, rather than a sequence of talks, compared to workshops or conferences. The seminar opened with a session of *lightning talks* that gave every participant two minutes to introduce themselves and mark their activities and interests in verification and in comparative evaluation in particular.

In order to give participants insight into different verification techniques and practical capabilities of some existing verification tools, the seminar featured short overviews of the state of the art in deductive verification resp. automatic verification, as well as several tutorials, hands-on sessions, and accompanying discussions. These included a longer tutorial on deductive verification with the Dafny system together with a hands-on session, as well as short mini-tutorials on the automatic verifiers CPAchecker and CBMC, and deductive verifiers KIV and VeriFast. Another hands-on session concluded the seminar.

Discussions on evaluation techniques and setups were initiated with presentations by competition organizers and benchmark collectors. The presented evaluation vehicles included: VerifyThis competition (deductive verification), SV-COMP (automatic verification), VSTTE competition (deductive verification), SMT-COMP (Satisfiability Modulo Theories), Run-time Verification competition, RERS challenge (Rigorous Examination of Reactive Systems), INTS benchmarks (Integer Numerical Transition Systems).

Since evaluation must be grounded with the requirements of current and prospective users of verification technology, the

seminar incorporated contributions from industrial participants. Among them were a talk on the use of the SPARK deductive verification tool-set as a central tool for the development of high-integrity systems at Altran UK, a talk on the use of automatic verification tools in the Linux Driver Verification project, an accompanying discussion, as well as statements by other industry representatives (from GrammaTech, Galois, LLNL, and Microsoft Research).

Verification Communities: Remarks on Commonalities, Differences, and Terminology

Important goals of the seminar were to raise mutual awareness and to foster cross-fertilization between verification communities. We may habitually refer to communities as "automatic verification" or "deductive verification", but as time passes, these labels become less adequate.

A trend is apparent that different types of tools are slowly converging, both technically and pragmatically. Instances of both automatic and deductive verifiers may use symbolic execution or SMT solvers. Automatic verifiers can synthesize (potentially quantified) invariants, verify infinite-state systems, or systems that are heap-centric.

The pace of development is high and the surveys are costly (the last comprehensive survey on automatic verification appeared in 2008). As a consequence, community outsiders typically have an outdated – sometimes by decades – view on verification technology that does not reflect the state of the art. We expect publications from competitions to fill the void between more comprehensive surveys.

One of the terminological pitfalls concerns the use of the attribute "automatic". For instance, model checking and data-flow analysis are typically advertised as "automatic". This is indeed true in the sense that the model-checking user does not have to supply proof hints such as loop invariants to the tool. On the other hand, using a model checker in practical situations may as well require user interaction for the purpose of creating test drivers, choosing parameters, tuning abstractions, or interpreting error paths (which can be quite complex). These aspects are typically abstracted away during evaluation of automatic verifiers, which allows better standardization but does not capture all aspects that are relevant in practice.

The situation is further confused by the fact that some deductive verifiers are also advertised as "automatic", even though all established deductive verification systems require user interaction and the amount of interaction that is needed with different tools is not radically different. The main meaningful differences are rather

1. whether user interaction happens only at the beginning of a single proof attempt or whether the user can/has to intervene during proof construction, and
2. whether user interaction happens in a purely textual manner or whether non-textual interaction is possible/required.

The seminar has confirmed the need for improved terminology, as well as made an attempt to eliminate misconceptions and communication pitfalls. Unfortunately, there is still no widely-accepted and usable terminology to communicate these distinctions.

4.27 Unifying Product and Software Configuration

Organizers: Krzysztof Czarnecki, Arnaud Hubaux, Ethan Jackson, Dietmar Jannach, and Tomi Männistö

Seminar No. 14172

Date: April 21–24, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.4.20

Participants: Michel Aldanondo, Danilo Beuche, Nikolaj Bjorner, Krzysztof Czarnecki, Conrad Drescher, Andreas Falkner, Jianmei Guo, Albert Haag, Holger H. Hoos, Lothar Hotz, Arnaud Hubaux, Dietmar Jannach, Christian Kästner, Eunsuk Kang, Charles Krueger, Wolfgang Küchlin, Daniel Le Berre, Tomi Männistö, Leonardo G. P. Murta, Klaus Schmid, Norbert Siegmund, Markus Stumptner, Juha Tiihonen, Andrzej Wasowski, Patrick Wischnewski, Ed Zulkoski

Customizable products are an integral part of most Business-to-Business (B2B) and Business-to-Consumer (B2C) markets. The fast-growing demand for mass-customization affects both tangible products (e. g., cars and mobile phones) and intangible products like software (e. g., operating systems, Enterprise Resource Planning systems and mobile phones). To this end, companies use software *configurators* that provide automated support to tailor products to the requirements of specific customers or market segments. These configurators have been developed essentially in two threads of research: *Product Configuration* (PC) and *Software Configuration* (SC).

PC is the umbrella activity of assembling and customizing physical artefacts (e. g., cars or muesli) or services (e. g., insurances). Due to the inherent complexity of configuration problems, PC was one of the first large-scale application fields of artificial intelligence (AI), as it required both powerful knowledge-representation formalisms and efficient reasoning methods. The particular challenges of knowledge representation and reasoning in PC even led to the development of new AI techniques. Today, PC can be seen as one of the major fields in which AI-based technology found its way into industrial practice and is part of many industrial configuration systems.

Mostly independent of PC, the software engineering community was confronted with challenging configuration problems. A typical challenge is the design and implementation of software components that can be adapted and parameterized according to customer requirements and business or technical constraints. As in PC approaches, the goal is to save costs by assembling individualized systems from reusable components. These challenges are dealt with in different strands of software engineering, e. g. software product line engineering or self-adaptive systems.

Questions of knowledge representation and types of reasoning support have been investigated for many years in PC and SC. Interestingly, research in these two fields has been carried out so far mostly independently. Except in rare cases, researchers in both fields are often unaware of approaches that have been developed in the other community.

This fragmentation is observable in two particular dimensions: *knowledge representation* and *configuration reasoning*. Knowledge representation is concerned with the question of how to encode the domain knowledge, e. g., about the compatibility of different features of a configurable product, in a formal or machine processible way. Configuration reasoning covers various aspects of how to make inferences given a knowledge base (configuration model), specific user requirements or an existing configuration. Typical tasks include the automatic completion of a partial configuration or checking the consistency of a given configuration.

The seminar was organized around the following research questions:

- (RQ1) What classes of configuration problems exist?
- (RQ2) How are these problems modelled?
- (RQ3) What automated tasks are supported?
- (RQ4) How are these automated tasks implemented?

The seminar was structured into three main blocks: *Problem characteristics*, *Knowledge representation* and *Reasoning and tools*. Each block consisted of a number of introductory presentations on the topic, which were given by researchers from different subfields and the seminar participants from industry. These talks then served as a basis for discussions on commonalities, differences and possible synergies. These discussions were made in small working groups in break-out sessions and the results then synthesized in plenary meetings. To make these break-out sessions more effective, the seminar participants were asked to fill out a detailed questionnaire before the seminar.

Overall, the seminar featured more than a dozen introductory talks from academia and from industry. In general, the interest

from industry was particularly encouraging and the seminar was attended by representatives and speakers, e. g., from IBM, SAP, Microsoft, Siemens and BigLever. The evening sessions were used by several seminar participants to give additional "lightning" talks, to share recent research results and dive deeper into technical aspects.

4.28 Multi-agent Systems and their Role in Future Energy Grids

Organizers: Michael N. Huhns, Wolfgang Ketter, Ryszard Kowalczyk, Fabrice Saffre, and Rainer Unland

Seminar No. 14181

Date: April 27 to May 2, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep4.4.37

Participants: Marco Aiello, Costin Badica, Tina Balke, Matthias Bürger, Liana Cipcigan, Hermann de Meer, Christian Derksen, Gilbert Fridgen, Hanno Hildmann, Michael N. Huhns, Micha Kahlen, Michael Kaisers, Stamatis Karnouskos, Wolfgang Ketter, Sonja Klingert, Matthias Klusch, Ryszard Kowalczyk, Winfried Lamersdorf, Sebastian Lehnhoff, Tobias Linnenberg, Christoph Mayer, Lars Mönch, Sascha Ossowski, Peter Palensky, Wolfgang Renz, Fabrice Saffre, Rainer Unland, Konstantina Valogianni, Marjan van den Akker, Eric van Heck, Minjie Zhang

Due to the depletion of scarce resources for energy production and the problems associated with climate change, there is widespread interest in new approaches formanaging energy generation, distribution, transportation, and consumption. Overall we must find a way to combine the economics, physics, physical components, and governmental policies and regulations of energy systems, while satisfying the personal preferences of consumers. The goal is to create a global Smart Grid.

The main differences between the current and the envisioned future grid are the production ecosystem on one hand and the information exchange on the other. Current grids traditionally rely on a comparatively stable number of large power plants that produce a constant and predictable amount of power, as well as on smaller power plants that can be activated quickly if demand requires it. Power and information flow from the supply side to the demand side. This is reflected in the underlying business models, which are mainly dictated by the prices the few big producers can achieve on the global market and by the costs for transmitting the power through a distribution system usually owned by private companies. This will change as more and more renewable and distributed generation technologies spread to a household level.

The distinction between producer and consumer will become increasingly blurred as the flow of power as well as information among the resultant *prosumers* becomes bi-directional. The current grid operates at a high-voltage level suited for long distance delivery, while a prosumer-based network will be a more localized and low-voltage grid. Further, the increasing use of renewable sources will result in a less predictable generation pattern, a matter which in itself is raising a number of interesting challenges. In short, the new power grids will differ in magnitude and direction as well as in generation consistency, which will require a complete revision of the underlying business model as the currently predominant global (or, at least, national) market will be replaced by a number of local markets that will have to maintain the balance between supply and (individually generated) demand, i. e., market places for power generation as well as power consumption.

The development of an information and communication technology (ICT) support infrastructure will be the key challenge in the design of an end-to-end smart grid framework. This will require the capability to balance supply and demand and to handle complex operations. The efficient, real-time exchange of information and the coordinated decisions among many stake holders (consumers, distributors, transporters, and generators) have to be supported. This is not possible within the structure and practice of the current grid. Different levels of the grid (layout, control, ICT infrastructure, maintenance, failure handling, and business models), as well as the communication and cooperation among these levels, needs to be fully coordinated with all the other levels. To predict the emergent properties of the system under a range of different conditions and worst-case scenarios, extensive and effective simulation tools will be required. A solution to this large and very complex problem requires intelligent decisions to be made at each node in the electric power distribution network (grid), especially at the edges. To be manageable, the decisions must take advantage of locality constraints and end-user preferences. A multiagent system, with agents located at the edges and nodes of the grid and representing the interests of end-users, distributors, and providers, satisfies these requirements. It is thus the default system solution thatwas considered first and adopted at the Dagstuhl.

4.29 Categorical Methods at the Crossroads

Organizers: Samson Abramsky, John Baez, Fabio Gadducci, and Viktor Winschel

Seminar No. 14182

Date: April 27 to May 2, 2014 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.4.4.49

Participants: Samson Abramsky, John C. Baez, Robin Cockett, Bob Coecke, Andrea Corradini, Andrée Ehresmann, Fabio Gadducci, Ulrike Golas, Michael Johnson, Bartek Klin, Barbara König, Tom S. Maibaum, Martin Meier, Francois Metayer, Ugo Montanari, Jeffrey C. Morton, Till Mossakowski, Fernando Orejas, Dusko Pavlovic, Andrew M. Pitts, Gordon Plotkin, Peter Selinger, Alexandra Silva, Alex Simpson, Pawel Sobocinski, Sam Staton, Jamie Vicary, Viktor Winschel, Philipp Zahn

Since the 1960s, category theory has been recognised as a powerful conceptual framework and a flexible specification language. The range of research areas where categorical methods found application is quite wide: from physics, economics, and linguistics to many branches of mathematics, especially algebraic geometry, algebraic topology, and logic. And, of course, computer science: possibly the discipline, apart from mathematics, where these methods have been most wholeheartedly adopted. Indeed, they have become part of the standard "tool-box" in many areas of theoretical informatics, from programming languages to automata, from process calculi to type theory.

Despite their flexibility and expressiveness, a more general acceptance of categorical methods has been hindered by the perceived difficulties of the formalism. As a consequence, many researchers in different communities share the feeling of under-exploitation of the potentialities of category theory to their areas of interest. This Dagstuhl Perspectives Workshop seminar brought together people from various disciplines that are interested in the application of categorical tools in their research area, from computer science towards venues such as economics, mathematics, and physics.

Besides the benefits to the understanding of each topic that came from a plurality of voices in a discussion, the meeting tried to address more general concerns. As far as some disciplines are concerned, the workshop helped the reconciliation of different research strands that uses categorical tools. Most importantly though, the workshop aimed at building bridges between disciplines, by reviewing the variety of uses of categorical methods in different fields and trying to find common abstractions that allow the same structures and concepts to be recognized as they arise in different settings, and to be transferred from one area to another.

In order to put on firm grounds the foundations of a common language, each working day included two survey lectures during the morning, which presented a variety of topics where categorical methods play a major role. These were followed by three shorter talks in the early afternoon, which presented active areas and innovative application for these methods. The day was closed by a working group session: during these sessions the attendants split into several groups according to the main thematic areas that had been identified on the first day. The variety of topics dealt with in the workshop was large, and the suggested application areas included (quantum) computation, physics, biology, complex systems, economic, social and cognitive science, and linguistics. Indeed, some of the items span more than one discipline, e. g. game theory, and the list is definitively not exhaustive.

Although the scope of the workshop was broad, the over-arching research theme was to develop categorical methods as a unified approach to the modeling of complex systems, and category theory as a paradigm for mathematical modeling and applied science. To this end, the overall purpose of the workshop was to start developing a coherent research community applying categorical methods to a wide range of disciplines. Under these terms, the workshop has been indeed successful. Laying out a common mathematical language and finding analogies among apparently distant concepts from unrelated disciplines provided a basis for fruitful cross-disciplinary interactions also by facilitating a "technology transfer". Concretely, this led to the fostering of new collaborations among the participants, and the preliminary exploration of new directions and research themes.

4.30 Horn Formulas, Directed Hypergraphs, Lattices and Closure Systems: Related Formalisms and Applications

Organizers: Kira Adaricheva, Giuseppe F. Italiano, Hans Kleine Büning, and György Turán

Seminar No. 14201

Date: May 11–16, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.5.1

Participants: Kira V. Adaricheva, Marta Arias, Giorgio Ausiello, Jose Luis Balcazar, Laurent Beaudou, Kristof Berczi, Karell Bertet, Endre Boros, Ondřej Čepek, Vincent Duquenne, Thomas Eiter, Donatella Firmani, John Franco, Loukas Georgiadis, Amélie Gheerbrant, Gabriel Istrate, Giuseppe F. Italiano, Roni Khardon, Hans Kleine Büning, Petr Kučera, Oliver Kullmann, Sergei O. Kuznetsov, Luigi Laura, Leonid Libkin, Kazuhisa Makino, Pierre Marquis, Umberto Nanni, Amedeo Napoli, James B. Nation, Lhouari Nourine, Sergei Obiedkov, Jean-Marc Petit, Sebastian Rudolph, Petr Savicky, Robert H. Sloan, Ewald Speckenmeyer, Despina Stasi, Suguru Tamaki, György Turán, Marcel Wild

Brief Introduction to the Topic

The Dagstuhl Seminar 14201 on "Horn formulas, directed hypergraphs, lattices and closure systems: related formalisms and applications" was motivated by the growing recognition in the respective research communities that theoretical research and applications of the areas would benefit from increasing the interaction between these fields of research.

These areas deal with very closely related concepts, but have traditionally been studied within logic, algebra, combinatorics, database theory and artificial intelligence using different techniques and often exploring similar questions with somewhat different emphasis corresponding to the particular area. One of the basic results, the existence of GD-basis, was discovered independently and has different proofs in several areas, such as database theory, the theory of implicational systems and computational learning theory.

The principal objective of the seminar was, as formulated in the seminar announcement, to "bring together a critical mass of researchers and to provide a platform for personal contacts and scientific interchange between the different disciplines in an atmosphere that will stimulate collaboration and lead to new partnerships". In particular, it was hoped that the invitation of a large number of young researchers who would then become familiar with the related research in all the topics discussed, will contribute to the fruitful study of these areas as a more unified discipline in the next generation. Another, related, objective was to help crystallize the main research directions and to disseminate challenging open problems across the different research areas.

Organization of the Seminar and Activities

The seminar brought together 40 participants working in various areas of mathematics and computer science, mostly in algebra, logic, date base theory, artificial intelligence and data mining. In order to establish the common ground for the discussion and use of terminology, the organizers planned five tutorial talks that were scheduled in the first two days of the seminar. There were the following:

- Endre Boros, in Horn Boolean functions
- Leonid Libkin, in Data Bases
- Marcel Wild, in Closure Systems
- Karell Bertet, in Implicational systems and Concept Analysis
- Georgio Ausiello, in Directed Hypergraphs

It is worth mentioning that two of the tutorial presenters, Endre Boros and Georgio Ausiello, are editors-in-chief of two leading journals that often publish papers associated with the topic of the seminar: Applied Discrete Mathematics and Theoretical Computer Science, respectively. Most other talks of the seminar were related to one or more of these big themes, and they were loosely grouped into sections of presentations in order to stimulate the discussion during regular sessions of the seminar, as well as the break time and follow-up informal meetings. The following grouping gives an approximation to the various topics reflected in the presentations:

- Closures: Duquenne, Wild, Rudolph, Khardon
- Implicational systems: Wild, Adaricheva, Bertet
- Databases: Petit, Libkin
- Directed hypergraphs: Berczi, Nanni, Turán, Ausiello
- Concept Analysis: Obiedkov, Napoli, Kuznetsov
- Horn formulas: Arias, Kučera, Stasi, Čepek
- Applications : Arias, Balcazar, Nation, Bertet
- Knowledge representation: Khardon, Marquis, Sloan
- Constraints: Tamaki

- Probability: Istrate, Balcazar
- Satisfiability: Kleine Büning, Kullmann, Wild

There was a relatively large group of participants attending the seminar without giving a presentation, many of them formed an active audience that initiated discussions and informal meetings. There were several young participants, who associated themselves with one or two presenters, some of them were co-authors in presented results. In some presentations, the participants touched the rare topics which remarkably found active response from the audience. For example, U. Nanni mentioned application of hypergraphs in the modeling of E-Learning. This seems to open a new connection to the theory Knowledge Spaces that were not previously considered by the organizers as the direction of common interests with the current seminar. S. Tamaki presented the overview of recent developments in constraint satisfiability, the topic of the separate Dagstuhl seminar in 2012, which is also familiar to a number of participants of the current seminar. This gave a nice connection to the earlier collective effort of researchers approaching a common research problem with different tools. On Wednesday most participants took part in the countryside walk, the venue of informal exchange of news and views. It was followed by a dinner and a session on open problems. 17 open problems were presented by the participants, some of them referring to the topics of presentations, several young participants introducing new topics. On Thursday night many informal groups were gathering in the after-dinner discussions. During the work of the whole seminar participants took advantage of open access to the comprehensive library, in particular, checking the collection of books authored by participants of the seminar.

Concluding Remarks

We believe that the seminar was very successful in bringing together a critical mass of researchers from different communities and in providing a platform for personal contacts and scientific interchange between the participants. Due to its highly interdisciplinary nature, in order to stimulate collaborations and to foster possible interactions between different research communities, the organizers decided to schedule talks and discussions not only grouped according to topics but also so as to provide a vivid mix of different research questions and results. In particular, the first two days started with introductory talks (tutorials and surveys) delivered by leading experts, while the rest of the seminar included talks presenting current research and applications as well as problem sessions aimed at identifying core research problems coming from the different fields. Besides presentations, the program offered room for open discussions and informal working groups. As a major outcome, a special issue of the journal *Theoretical Computer Science*, co-edited by the organizers, will be devoted to the themes of the seminar. We hope this could serve as a reference material for future interdisciplinary research in the field. Schloss Dagstuhl and its staff provided a very convenient and stimulating environment. The seminar participants appreciated the cordial atmosphere which improved mutual understanding and inspiration. The organizers of this seminar wish to thank all those who helped to make the seminar a fruitful research experience.

4.31 JA4AI – Judgment Aggregation for Artificial Intelligence

Organizers: Franz Dietrich, Ulle Endriss, Davide Grossi, Gabriella Pigozzi, and Marija Slavkovik

Seminar No. 14202

Date: May 12–15, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.5.27

Participants: Edmond Awad, Dorothea Baumeister, Veronica Block, Franz Dietrich, Ulle Endriss, Umberto Grandi, Davide Grossi, Paul Harrenstein, Frederik S. Herzberg, Ron Holzman, Christian Klamler, Sébastien Konieczny, Lewis A. Kornhauser, Jérôme Lang, Yixi Li, Christian List, Xavier Mora, Ilan Nehama, Klaus Nehring, Nicolas Paget, Gabriella Pigozzi, Ramon Pino-Pérez, Marcus Pivato, Daniele Porello, Clemens Puppe, Marija Slavkovik, Paolo Turrini, Leon van der Torre, William S. Zwicker

Judgment aggregation is a group decision-making theory, developed in the last decade, that studies how to reach group decisions on logically interconnected issues by aggregation of individual decisions on those issues. The interest of computer science in group reasoning and decision-making theories is driven by the increase of distribution of information and computation as features of various Internet-based services that dominate the information technology market.

Judgment aggregation studies collective decision-making as a process whereby individual opinions concerning the acceptance or rejection of a set of issues are aggregated into one collective judgment. The problem is for the aggregation process to preserve, in a non-trivial way, some 'rational' aspects of the individual to-be-aggregated stances like, in particular, logical consistency. A wealth of results have highlighted how the rationality of a collective decision may clash with other desirable properties of a process of aggregation one may wish to require (*e.g.*, anonymity of the voters, independence of the aggregated issues, to mention a few).

Judgment aggregation research, originally studied in law, was propelled into other disciplines with its establishment as a separate discipline from preference aggregation in the early 2000's. The first half of the decade was marked by studies of aggregation properties that cannot be jointly satisfied by one aggregation function, usually referred to as 'impossibility results'. These studies were mostly conducted by researchers from political science, law, economics, mathematics, and philosophy. The second half of the decade witnessed an increase of interest in judgment aggregation of researchers from artificial intelligence (AI), specifically knowledge representation and reasoning (KR), and multi-agent systems (MAS).

Research on judgment aggregation, from the computer scientific perspective, has splintered in many directions, with scholars pursuing very different lines of research: judgment aggregation and logic, judgment aggregation and complexity theory, judgment aggregation and relations to preference aggregation, judgment aggregation and belief merging, judgment aggregation and argumentation, to mention a few. At the same time work in judgment aggregation has diversified in non-computer science disciplines: judgment aggregation and deliberation, judgment aggregation and strategic voting, judgment aggregation and probabilistic opinion pooling, to mention a few. Despite the common research thread, having so many disciplines involved make it difficult to keep track of the research advancements across all domains.

The goal of this Dagstuhl seminar was to give researchers across the contributing disciplines an integrated overview of the current research and interests in judgment aggregation and of its emerging trends, and by doing this, to kick-start a lasting interdisciplinary network bridging the computer science/humanities divide in the field. To accomplish this goal, we structured the seminar around four types of events:

- Invited tutorials – three invited overview talks aimed to introduce the interdisciplinary audience to the origins and advancements of judgment aggregation in law, political science and computer science.
- Contributed talks – fourteen contributed talks of thirty minutes each.
- Networking sessions – two free networking sessions.
- Rump session – open to all participants to present new ideas.

The topics of the invited talks were chosen so as to give a foundation of the disciplines in which judgment aggregation originated and was formalised, as well as to motivate the interest of judgment aggregation for computer science. Although we expected that all of the participants would be familiar with at least one of these foundational topics, we also expected them to be unfamiliar with at least one as well. The tutorial lectures aimed to

homogenise the background knowledge in judgment aggregation among the participants.

The contributed talks aimed to introduce the community with the recent work of the speakers. We accommodated fourteen talks, possibly compromising on the length of the talk itself in the interest of allowing space for questions. We are happy to observe that there was a lively debate after each of the talks, which we expect shall contribute towards advancement of each of the presented works.

Given the short period of three days and prior Dagstuhl experience of the organisers, we decided to not structure the networking session and simply allow for a time for the participants to talk to each other and get to know about each other's work and interests. The enthusiastic discussions following the contributed talks typically continued into the networking sessions.

The rump session was free for a last-minute sign up to all participants. Each interested person was given a five-minute time slot to present an idea that emerged during the seminar or a work in progress. A third of the participants took this opportunity to present. This was a very lively and well received part of the seminar. In retrospect, a similar session would have been well received also at the beginning of the seminar, giving the participants more time to discuss the presented ideas.

4.32 The Future of Refactoring

Organizers: Danny Dig, William G. Griswold, Emerson Murphy-Hill, and Max Schäfer

Seminar No. 14211

Date: May 18–23, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.5.40

Participants: Don Batory, Ira D. Baxter, Andrew P. Black, Robert Bowdidge, John Brant, Caius Brindescu, Marcio Cornelio, Stephan Diehl, Danny Dig, Ran Ettinger, Alejandra Garrido, Rohit Gheyi, William G. Griswold, Shinpei Hayashi, Felienne Hermans, Jongwook Kim, Huiqing Li, Francesco Logozzo, Kim Mens, Tom Mens, Naouel Moha, Emerson Murphy-Hill, Oscar M. Nierstrasz, Bill Opdyke, Chris Parnin, Javier Perez, Veselin Raychev, Don Roberts, Max Schaefer, Gustavo Soares, Peter Sommerlad, Friedrich Steimann, Kathryn T. Stolee, Volker Stolz, Simon J. Thompson, Eli Tilevich, Frank Tip, Mohsen Vakilian, Jurgen Vinju, Louis Wasserman, Jan Wloka

The Dagstuhl seminar on "The Future of Refactoring" brought together 41 researchers and practitioners from academia and industry working on different aspects of refactoring. Participants had the opportunity to introduce their own work both in short plenary talks and more detailed presentations during breakout sessions, with daily keynote talks by eminent refactoring researchers providing historical background. Given the rapid growth of the field over the past decade, special emphasis was put on providing opportunities for researchers with similar interests to meet and survey the state of the art, identify open problems and research opportunities, and jointly chart the future of refactoring research.

We believe the seminar achieved its goal of providing a forum for in-depth discussion of recent research in the area, and of fostering collaboration. In particular, it kickstarted several collaborative projects, among them a book on refactoring tools, a special journal issue on refactoring and a survey article on refactoring research over the last decade.

Research Context

Modern software is rarely written from scratch. Instead, it usually incorporates code from previous systems, and is itself reincarnated in other programs. Modern software is also not static. Instead, it constantly changes as bugs are fixed and features added, and usually these changes are performed by more than one programmer, and not necessarily by the original authors of the code.

Refactoring is a technique for supporting this highly dynamic software life cycle. At its most basic, refactoring is the process of improving the internal structure of an existing piece of code without altering its external behavior. It can be used for cleaning up legacy code, for program understanding, and as a preparation for bug fixes or for adding new features. While any behavior-preserving change to a program can be considered a refactoring, many particularly useful and frequently recurring refactoring operations have been identified and catalogued. Over the past decade, popular development environments have started providing automated support for performing common refactorings, making the process of refactoring less tedious and error-prone.

Based on the accumulated experience with refactorings both in practical applications and in research, this seminar aimed to identify open problems and challenges and to foster collaboration between researchers and between academia and industry to address these issues and actively shape the future of refactoring.

Seminar Format

Given the large number of participants, the standard conference format with one in-depth talk per participant would have been impractical. Instead, we decided to split up the schedule: during the first three days, the mornings were allocated to plenary sessions. Each day began with a keynote by a distinguished speaker with decades of experience with refactoring, in which they presented their perspective on refactoring. The rest of the morning was allocated to "lightning talks" where each participant was given a 7-minute presentation slot for providing a quick, high-level overview of their work without getting bogged down in detail, followed by a few minutes for questions. While this format was not easy for the speakers, everyone rose to the challenge, and reactions from both presenters and audience were broadly positive.

Monday afternoon was given over to four parallel breakout sessions organized along thematic lines: novel domains for refactoring, user experience in refactoring, refactoring tools and meta-tools, and refactoring in education. While participants appreciated the opportunity for more in-depth presentations and discussion, this format had the unfortunate but inevitable draw-

back that several talks were held in parallel, and not everyone was able to attend all the talks they were interested in.

Tuesday afternoon had an industry panel, followed by another round of breakout sessions. Discussion and exchange continued in an informal setting during Wednesday afternoon's excursion to Mettlach.

On Thursday morning, we had another keynote followed by a final round of breakout sessions. While the focus of the breakout sessions on Monday and Tuesday had been on surveying recent work and getting an overview of the state of the art, Thursday's sessions were aimed at gathering together the threads, and identifying common themes, open problems and research opportunities.

The outcome of these group discussions were then briefly presented in a plenary on Thursday afternoon, and opportunities for collaborative projects were identified. Specifically, the following projects were discussed and planned in group discussions on Thursday afternoon:

- a book on refactoring tools;
- a special issue of IEEE Software on refactoring;
- a survey paper on refactoring research in the last decade;
- an informal working group on the place of refactoring in the Computer Science curriculum.

Friday morning saw a final plenary discussion, summarizing the project discussions of Thursday afternoon and ending with a retrospective session on which aspects of the seminar are worth keeping for the future, what needs to change, and what still puzzles us.

We hired George Platts, a professional artist, to facilitate games he designed and tangential thinking activities to help the participants develop a sense of scientific community. During each of the five days of the Seminar, George ran 30-minute games sessions at the beginning of the day which doubled as times for announcements to be given and daily reports to be delivered. In the early afternoon, we had a 30-minute game session to energize participants for the afternoon's workshops. For the rest of the time in his 'studio', he has been playing music, showing short films, facilitating drawing and painting activities, composing sound composition for all participants to perform.

4.33 Geometric Modeling

Organizers: Falai Chen, Tor Dokken, Thomas A. Grandine and Stefanie Hahmann

Seminar No. 14221

Date: May 25–30, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.5.68

Participants: Oliver Barrowclough, Pere Brunet, Laurent Busé, Falai Chen, Elaine Cohen, Tor Dokken, Gershon Elber, Ioannis Z. Emiris, Ron Goldman, Thomas A. Grandine, Jens Gravesen, Hans Hagen, Stefanie Hahmann, Kai Hormann, Bert Jüttler, Myung-Soo Kim, Tae-Wan Kim, Leif Kobbelt, Rimvydas Krasauskas, Nicole Lehmann, Yaron Lipman, Ligang Liu, Tom Lyche, Geraldine Morin, Bernard Mourrain, Georg Muntingh, Peter Noertoft, Jörg Peters, Konrad Polthier, Helmut Pottmann, Hartmut Prautzsch, Ewald Quak, Ulrich Reif, Richard F. Riesenfeld, Jarek Rossignac, Malcolm A. Sabin, Maria Lucia Sampoli, Scott Schaefer, Vadim Shapiro, Hiromasa Suzuki, Georg Umlauf, Tamas Varady, Nelly Villamizar, Wenping Wang

The 9th Dagstuhl seminar was attended by 44 leading researchers coming from 3 continents and 20 countries. A total of 45 presentations were grouped together into 12 lecture sessions and 3 perspective working group sessions. There was also ample time for stimulating and fruitful person to person and group discussions in the harmonic Dagstuhl atmosphere. Dagstuhl seminars on Geometric Modelling are among the most interdisciplinary events within Geometric Modelling. The reason is the seminar format and the generous numbers of targeted invitation to leading researchers across the different research communities addressing Geometric Modelling. Geometric Modeling is the branch of Computer Science concerned with the acquisition, representation, modeling and analysis of 3-dimensional and higher dimensional geometry. The evolution of IT-technology with Cloud Computing and the big data challenge, and novel manufacturing technologies such as 3D printing and layered manufacturing, as well as the introduction of Isogeometric Analysis drive a need for increased innovation within Geometric Modeling. The Dagstuhl seminars on Geometric Modelling are one of the main driving forces facilitating such innovation.

The twelve lecture sessions covered a wide range of topics:

- Geometric modelling, analysis and computations;
- Methods in approximate algebraic geometry and implicitization;
- Mesh processing both related to triangulations and isogeometric analysis;
- Optimization and Isogeometric Analysis;
- Splines over triangulations an and locally refinable splines;
- Material modelling and reverse engineering;
- Funding opportunities for Geometric Modelling within Horizon 2020.

The three perspective working groups have each written a short document reporting on approach and result of the session. The perspective working group on subdivision addressed the state of the art and the future research challenges of subdivisions. The group on 3D printing approached the challenges from the current wide media coverage of 3D printing and challenges experienced by industry using 3D printing. The group on modeling of material microstructures was approaching the topic from additive manufacturing, and the challenges faced by the geometric modelling community to support modeling and representation of such material structures in variety of applications. As with all previous Dagstuhl Seminars on Geometric Modeling, the conference proceedings will be published, the 2014 proceedings as a special issue of the Elsevier Journal Graphical Models. A special event during the conference was the John Gregory Memorial Award honoring Elaine Cohen, Jörg Peters and Ulrich Reif. This award is presented every three years at Dagstuhl and honors fundamental contributions to the field of geometric modeling. The organizers thank all the attendees for their contributions and extend special thanks to the team of Schloss Dagstuhl for helping to make this seminar a success. As always, we enjoyed the warm atmosphere of the Schloss, which supports both formal presentations as well as informal exchanges of ideas.

Fig. 4.8
Drawing for the Dagstuhl children's guest book by Oscar (3), son of Dagstuhl Seminar 14132 participants Maike and Kevin Buchin.

4.34 Scientific Visualization

Organizers: Min Chen, Charles D. Hansen, Penny Rheingans, and Gerik Scheuermann

Seminar No. 14231

Date: June 1–6, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.6.1

Participants: James Ahrens, Georges-Pierre Bonneau, Rita Borgo, Nadia Boukhelifa, Hamish Carr, Remco Chang, Jian Chen, Min Chen, Hank Childs, João Luiz Dihl Comba, Leila De Floriani, Thomas Ertl, Issei Fujishiro, Christoph Garth, Andreas Gerndt, Michael Gleicher, Eduard Gröller, Hans Hagen, Charles D. Hansen, Helwig Hauser, Hans-Christian Hege, Mario Hlawitschka, Ingrid Hotz, Christopher R. Johnson, Kenneth Joy, Robert Michael Kirby, David H. Laidlaw, Heike Leitte, Ross Maciejewski, Georgeta Elisabeta Marai, Torsten Möller, Kristi Potter, Penny Rheingans, Timo Ropinski, Gerik Scheuermann, Claudio T. Silva, Holger Theisel, Amitabh Varshney, Anna Vilanova Bartroli, Daniel Weiskopf, Rüdiger Westermann, Ross Whitaker, Anders Ynnerman, Xiaoru Yuan

Scientific Visualization (SV) is the transformation of digital data, derived from observation or simulation, into readily comprehensible images, and has proven to play an indispensable part of the scientific discovery process in many fields of contemporary science. Since its inception two decades ago, the techniques of Scientific Visualization have aided scientists, engineers, medical practitioners, and others in the study of a wide variety of data including, for example, high-performance computing simulations, measured data from scanners (CT, MR, confocal microscopy, satellites), internet traffic, and financial records. One of the important themes being nurtured under the aegis of Scientific Visualization is the utilization of the broad bandwidth of the human sensory system in steering and interpreting complex processes and simulations involving voluminous data across diverse scientific disciplines. Since vision dominates our sensory input, strong efforts have been made to bring the mathematical abstraction and modeling to our eyes through the mediation of computer graphics. This interplay between various application areas and their specific problem-solving visualization techniques has been the goal of all the Dagstuhl Scientific Visualization seminars and was emphasized in the seminar which took place June 1–6, 2014.

Our seminar was focused on four research themes that will have significant impact in the coming years. These four themes reflect the heterogeneous structure of Scientific Visualization and the current unsolved problems in the field. They represent cross-cutting topic areas where applications influence basic research questions on one hand while basic research drives applications on the other. This cross-cutting feature makes Dagstuhl a unique setting in the research community, as the scientific coverage of the seminar is broader than other more focused workshops and seminars hosted at Dagstuhl while much more focused and forward-looking than general conferences. Our four themes were:

Uncertainty Visualization: Decision making, especially rapid decision making, is always made under uncertain conditions. As former English Statesman and Nobel Laureate (Literature), Winston Churchill said, "True genius resides in the capacity for evaluation of uncertain, hazardous, and conflicting information." and echoed by Nobel Prize winning physicist Richard Feynman, "What is not surrounded by uncertainty cannot be the truth." Uncertainty visualization seeks to provide a visual representation of errors and uncertainty for three-dimensional visualizations. Challenges include the inherent difficulty in defining, characterizing, and controlling comparisons among different data sets and in part to the corresponding error and uncertainty in the experimental, simulation, and/or visualization processes.

Integrated Multi-field Visualization: The output of the majority of computational science and engineering simulations is typically a combination of fields, generally called multi-field data, involving a number of scalar fields, vector fields, or tensor fields. Similarly, data collected experimentally is often multi-field in nature (and from multiple sources). The ability to effectively visualize multiple fields simultaneously, for both computational and experimental data, can greatly enhance scientific analysis and understanding. Multi-scale problems with scale differences of several orders of magnitude in computational fluid dynamics, material science, nanotechnology, biomedical engineering and proteomics pose challenging problems for data analysis. The state of the art in multi-scale visualization considerably lags behind that of multi-scale simulation. Novel solutions to multi-scale and multi-field visualization problems have the potential for a large impact on scientific endeavors.

Environmental Scientific Visualization: Environmental scientific visualization or environmental visualization refers to a collection of visualization applications that deal with cap-

tured and simulated data in climate research, atmospheric and environmental sciences, earth science, geophysics and seismic research, oceanography, and the energy industry (e. g., oil, gas and renewable energy). Research in these application domains has a huge impact on mankind, and typically faces serious challenges of data deluge (e. g., very large volumes of multi-spectral satellite images, large data collections from different sensor types, ensemble computation of very large simulation models, scattered, time-varying, multi-modal data in seismic research). In comparison with biomedical visualization and small-to-medium scale computational fluid dynamics, the effort for developing visualization techniques for such applications has not been compatible with the importance and scale of the underlying scientific activities in these application domains. Scientific progress in the areas of the environment and sustainability is critical in the solution of global problems and scientific visualization has great potential to support this progress.

Scientific Foundation of Visualization: The rapid advances in scientific visualization have resulted in a large collection of visual designs (e. g., for flow visualization), algorithms (e. g., for volume rendering), and software tools and development kits. There have also been some scattered investigations into the theoretic and perceptual aspects of visualization. However, many fundamental questions remain unanswered, such as, why is one visual design more effective than another, can visual designs be optimized and how, what is the role of visualization in a scientific workflow and how can such a role be formalized in a scientific workflow, can visualization quality be measured quantitatively and how, and what is the most effective way to conduct perceptual and usability studies involving domain experts? With the experience of delivering technical advances over the past two decades, it is timely for the visualization community to address these fundamental questions with a consorted effort. Such an effort will be critical to the long-term development of the subject, especially in building a scientific foundation for the subject.

The format of the seminar was two-part: having groups of four to five shorter talks followed by a panel of the speakers which encouraged discussion and breakout groups on the four topics as well as topics which came up at the meeting. The scientific presentations were scheduled at the beginning of the week in order to simulate the discussions from a broad perceptive. Unlike the typical arrangement, all presentations in each session were given in sequence without a short Q&A session at the end of each talk. Instead, all speakers of a session were invited to sit on the stage after the presentation, and answer questions in a manner similar to panel discussions. This format successfully brought senior and junior researchers onto the same platform, and enabled researchers to seek a generic and deep understanding through their questions and answers. It also stimulated very long, intense, and fruitful discussions that were embraced by all participants. The breakout groups focused on the general themes and are reported in the following sections.

4.35 Design and Synthesis from Components

Organizers: Jakob Rehof and Moshe Y. Vardi

Seminar No. 14232

Date: June 1–6, 201401 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.6.29

Participants: Farhad Arbab, Christel Baier, Ugo de'Liguoro, Mariangiola Dezani, Laurent Doyen, Boris Düdder, Bernd Finkbeiner, Dror Fried, George T. Heineman, Fritz Henglein, Paola Inverardi, Neil D. Jones, Bengt Jonsson, Axel Legay, Moritz Martens, Hongyang Qu, Jean-François Raskin, Jakob Rehof, Sven Schewe, Joseph Sifakis, Bernhard Steffen, Wolfgang Thomas, Massimo Tivoli, Stavros Tripakis, PawełUrzyczyn, Moshe Y. Vardi, Clemens Wiltsche, Eran Yahav

The purpose of the seminar was bringing together researchers from the component-oriented design community, researchers working on interface theories, and researchers working in synthesis, in order to explore the use of component- and interface design in program synthesis.

The seminar proposal was motivated by a recently developing trend in component-based synthesis, which is seen both as creating a need and providing the potential for a cross-community effort. Traditionally, synthesis has been pursued in two distinct and somewhat independent technical approaches. In one approach, synthesis is characterized by temporal logic and automata theoretic methods, whereas in the other synthesis is characterized by deductive methods in program logics and in type theory considered under the Curry-Howard isomorphism. Recent work in component-oriented design has spurred the idea of *component-based synthesis*, where systems are synthesized relative to a given collection (library, repository) of components, within both technical approaches. Recent results in both communities show that this development allows the two communities to communicate more intensely on the common ground of component-orientation to their mutual benefit. The trend opens the door to a new attack on the great challenges of synthesis (including computational complexity and complexity of specification) by exploiting component design.

The seminar program consisted of 6 tutorial talks (1 hour) and 16 contributed talks (30 mins) as well as joint discussion sessions. Two slots for joint discussions were pre-planned for each day but were used flexibly and dynamically, depending on the development of discussions and reactions to the talks. It was felt that the mixture of tutorials, talks and joint discussion slots turned out to be an altogether very good instrument for making intensive exchanges among all seminar participants possible. It seems to be the general impression that the seminar was very succesful in meeting the challenge of bringing together researchers from quite a diverse range of technical fields, spanning from software engineering to mathematical logic. The seminar was succesful in generating several concrete cross-community collaboration projects which would not have been likely to have come into existence by way of traditional conferences.

Joint discussions were summarized by Dror Fried (Rice University) who is gratefully acknowledged for undertaking the role of "seminar collector".

4.36 Challenges in Analysing Executables: Scalability, Self-Modifying Code and Synergy

Organizers: Roberto Giacobazzi, Axel Simon, and Sarah Zennou

Seminar No. 14241

Date: June 9–13, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.6.48

Participants: Davide Balzarotti, Sébastien Bardin, Frederic Besson, Sandrine Blazy, Juan Caballero, Lorenzo Cavallaro, Aziem Chawdhary, Cory Cohen, Mila Dalla Preda, Bjorn De Sutter, Saumya K. Debray, David Delmas, Thomas Dullien, Emmanuel Fleury, Anthony Fox, Roberto Giacobazzi, Kathryn E. Gray, Paul Irofti, Yan Ivnitskiy, Andy M. King, Tim Kornau, Julian Kranz, Colas Le Guernic, Junghee Lim, Alexey Loginov, Federico Maggi, Jean-Yves Marion, Florian Martin, Isabella Mastroeni, Bogdan Mihaila, Magnus Myreen, Gerald Point, Edward Robbins, Bastian Schlich, Alexander Sepp, Axel Simon, Aditya Thakur, Axel Tillequin, Franck Védrine, Aymeric Vincent, Xueguang Wu, Brecht Wyseur, Stefano Zanero, Sarah Zennou

As a follow-up on the previous Dagstuhl Seminar 12051 on the analysis of binaries, the interest in attending this new seminar was very high. In the end, less than half the people that we considered inviting could attend, namely 44 people. In contrast to the previous seminar that ran for 5 days, this seminar was a four-day seminar due to a bank holiday Monday. Having arranged the talks by topic, these four days split into two days on the analysis of binaries and into (nearly) two days on obfuscation techniques.

The challenges in the realm of general binary analysis have not changed considerably since the last gathering. However, new analysis ideas and new technologies (e. g. SMT solving) continuously advance the state-of-the-art and the presentations where a reflection thereon. With an even greater participation of people from industry, the participants could enjoy a broader view of the problems and opportunities that occur in practice. Given the tight focus on binary code (rather than e. g. Java byte code), a more detailed and informed discussion ensued. Indeed, the different groups seem to focus less on promoting their own tools rather than seeking collaboration and an exchange of experiences and approaches. In this light, the seminar met its ambition on synergy. It became clear that creating synergy by combining various tools is nothing that can be achieved in the context of a Dagstuhl Seminar. However, the collaborative mood and the interaction between various groups give hope that this will be a follow-on effect.

The second strand that crystallized during the seminar was the practical and theoretic interest in code obfuscation. Here, malware creators and analysts play an ongoing cat-and-mouse game. A theoretic understanding of the impossibility of winning the game in favor of the analysts helps the search for analyses that are effective on present-day obfuscations. In practice, a full understanding of some obfuscated code may be unobtainable, but a classification is still possible and useful. The variety of possible obfuscations creates many orthogonal directions of research. Indeed, it was suggested to hold a Dagstuhl Seminar on the sole topic of obfuscation.

One tangible outcome of the previous Dagstuhl Seminar is our GDSL toolkit that was presented by Julian Kranz. We believe that other collaborations will ensue from this Dagstuhl Seminar, as the feedback was again very positive and many and long discussions where held in the beautiful surroundings of the Dagstuhl grounds. The following abstracts therefore do not reflect on the community feeling that this seminar created. Please note that not all people who presented have submitted their abstracts due to the sensitive nature of the content and/or the organization that the participants work for.

4.37 Software Development Analytics

Organizers: Harald Gall, Tim Menzies, Laurie Williams, and Thomas Zimmermann

Seminar No. 14261

Date: June 22–27, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.6.64

Participants: Bram Adams, Alberto Bacchelli, Ayse Bener, Trevor Carnahan, Serge Demeyer, Premkumar T. Devanbu, Stephan Diehl, Michael W. Godfrey, Alessandra Gorla, Georgios Gousios, Mark Grechanik, Michaela Greiler, Abram Hindle, Reid Holmes, Miryung Kim, Andrew J. Ko, Lucas M. Layman, Andrian Marcus, Nenad Medvidovic, Tim Menzies, Leandro L. Minku, Audris Mockus, Brendan Murphy, Meiyappan Nagappan, Alessandro Orso, Martin Pinzger, Denys Poshyvanyk, Venkatesh-Prasad Ranganath, Romain Robbes, Martin Robillard, Guenther Ruhe, Per Runeson, Anita Sarma, Emad Shihab, Diomidis Spinellis, Margaret-Anne Storey, Burak Turhan, Stefan Wagner, Patrick Wagstrom, Jim Whitehead, Laurie Williams, Dongmei Zhang, Thomas Zimmermann

Software and its development generate an inordinate amount of data. For example, check-ins, work items, bug reports and test executions are recorded in software repositories such as CVS, Subversion, GIT, and Bugzilla. Telemetry data, run-time traces, and log files reflect how customers experience software, which includes application and feature usage and exposes performance and reliability. The sheer amount is truly impressive:

- As of July 2013, Mozilla Firefox had 900,000 bug reports, and platforms such as Sourceforge.net and GitHub hosted millions of projects with millions of users.
- Industrial projects have many sources of data at similar scale.

But how can this data be used to improve software? Software analytics takes this data and turns it into actionable insight to inform better decisions related to software. Analytics is commonly used in many businesses—notably in marketing, to better reach and understand customers. The application of analytics to software data is becoming more popular.

To a large extent, software analytics is about what we can learn and share about software. The data include our own projects but also the software projects by others. Looking back at decades of research in empirical software engineering and mining software repositories, software analytics lets us share all of the following:

- **Sharing insights.** Specific lessons learned or empirical findings. An example is that in Windows Vista it was possible to build high-quality software using distributed teams if the management is structured around code functionality (Christian Bird and his colleagues).
- **Sharing models.** One of the early models was proposed by Fumio Akiyama and says that we should expect over a dozen bugs per 1,000 lines of code. In addition to defect models, plenty of other models (for example effort estimation, retention and engagement) can be built for software.
- **Sharing methods.** Empirical findings such as insights and models are often context-specific, e. g., depend on the project that was studied. However, the method ("recipe") to create findings can often be applied across projects. We refer to *"methods"* as the techniques by which we can transform data into insight and models.
- **Sharing data.** By sharing data, we can use and evolve methods to create better insight and models.

The goal of this seminar was to build a roadmap for future work in this area. Despite many achievements, there are several challenges ahead for software analytics:

- How can we make data useful to a wide audience, not just to developers but to anyone involved in software?
- What can we learn from the vast amount of unexplored data?
- How can we learn from incomplete or biased data?
- How can we better tie usage analytics to development analytics?
- When and what lessons can we take from one project and apply to another?
- How can we establish smart data science as a discipline in software engineering practice and research as well as education?

Seminar Format

In this seminar, we brought together researchers and practitioners from academia and industry who are interested in empirical software engineering and mining software repositories to share their insights, models, methods, and/or data. Before the seminar, we collected input from the participants through an online survey to collect relevant themes and papers for the seminar. Most themes from the survey fell into the categories of method (e. g., measurement, visualization, combination of qualitative with quantitative methods), data (e. g., usage/telemetry,

security, code, people, etc.), and best practices and fallacies (e. g., how to choose techniques, how to deal with noise and missing data, correlation vs. causation). A theme that also emerged in the pre-Dagstuhl survey was analytics for the purpose of theory format, i. e., *"data analysis to support software engineering theory formation (or, data analytics in support of software science, as opposed to software engineering)"*.

At the seminar, we required that attendees

1. discuss the next generation of software analytics;
2. contribute to a *Software Analytics Manifesto* that describes the extent to which software data can be exploited to support decisions related to development and usage of software.

Attendees were required to outline a set of challenges for analytics on software data, which will help to focus the research effort in this field. The seminar provided ample opportunities for discussion between attendees and also provide a platform for collaboration between attendees since our time was divided equally between:

1. Plenary sessions where everyone gave short (10 minute) presentations on their work.
2. Breakout sessions where focus groups worked on shared tasks.

Our schedule was very dynamic. Each day ended with a "think-pair-share" session where some focus for the next day was debated first in pairs, then shared with the whole group. Each night, the seminar organizers would take away the cards generated in the "think-pair-share" sessions and use that feedback to reflect on how to adjust the next day's effort.

4.38 Scripting Languages and Frameworks: Analysis and Verification

Organizers: Fritz Henglein, Ranjit Jhala, Shriram Krishnamurthi, and Peter Thiemann

Seminar No. 14271

Date: July 1–4, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.6.84

Participants: Achim D. Brucker, Niels Bjoern Bugge Grathwohl, Ravi Chugh, Arlen Cox, Christos Dimoulas, Julian Dolby, Matthias Felleisen, Daniele Filaretti, Cormac Flanagan, Jeffrey Foster, Ronald Garcia, Philippa Gardner, Michael Greenberg, Arjun Guha, Shu-Yu Guo, Christian Hammer, Fritz Henglein, Roberto Ierusalimschy, Thomas Jensen, Ranjit Jhala, Matthias Keil, Shriram Krishnamurthi, Benjamin Lerner, Benjamin Livshits, Sergio Maffeis, Matt Might, Yasuhiko Minamide, Anders Møller, Joe Gibbs Politz, Ulrik Terp Rasmussen, Tamara Rezk, Tiark Rompf, Sukyoung Ryu, Alan Schmitt, Jeremy G. Siek, Gareth Smith, Manu Sridharan, Éric Tanter, Peter Thiemann, Sam Tobin-Hochstadt, Tom Van Cutsem, David Van Horn, Panagiotis Vekris, Ben Wiedermann, Kwangkeun Yi

In the past decade scripting languages have become more mature: the wild experimentation and almost wilful embrace of obfuscation by Perl has been replaced by the level-headed simplicity of Python and the embrace of programming language research roots by Ruby. As a result, these languages have moved into the mainstream: every Web user relies on JavaScript.

The Challenges of Scripting Languages

Though scripting languages have become more mature, from the perspective of building robust, reliable software, they still suffer from several distinct problems, each of which creates new challenges for the research community.

- While these languages have textual definitions, they lack more formal descriptions, and in practice the textual "definitions" are themselves often in conflict with the normative nature of the implementations. This is in contrast to languages like Standard ML where the formal definition comes first. *How far can we go in creating formal semantics from a combination of implementations and textual documents?*
- Tests – more than either implementations, textual definitions, or formal semantics – are becoming the norm for specification. For instance, the latest JavaScript standard explicitly embraces testing by publishing and regularly updating a conformance suite. Similarly, a team trying to create an alternate implementation of one of these languages may read the definition but what they really aspire to match is the test suite behavior. *How can we support test suites as a new avenue of programming language specification?*
- One of the reasons programmers find these languages enjoyable (initially) is that they offer a variety of "convenient" features, such as overloading. As programs grow, however, understanding the full – and unintended! – behaviors of programs becomes a non-trivial effort. *How can we design semantics and static and dynamic tools that can cope with the heavily understated and overloaded behaviors that make scripting languages attractive?*
- Programmers increasingly do not program in languages but in high-level frameworks built atop them. For instance, though "Ruby" is popular for Web programming, programmers rarely write Web applications directly in Ruby, but rather atop the higher-level Ruby on Rails platform. The result of imposing significantly higher-level interfaces is that they necessitate new reasoning modes. For instance, while the jQuery library is a pure JavaScript program, type-checking jQuery as if it were "merely" JavaScript would produce types that are both unreadably compex and relatively useless. *Can we build custom reasoning at the level of the frameworks, then we can provide views of these frameworks that are consistent with the level at which developers think of them, and can we check that the implementations adhere to these interfaces?*
- These languages and frameworks are themselves not enough. They all reside in an eco-system of a family of other languages and frameworks whose interdependencies are necessary for proper understanding of program execution. For instance, in the client-side Web, JavaScript – which has gotten significant attention from the research community – only runs in response to stimuli, which are obtained from the DOM. In turn, the DOM and JavaScript both depend on the style-sheets written in CSS. But in fact all three of these components – the JavaScript code, the CSS styling, and the DOM events – all depend on one another, because almost any one can trigger or modify the other. *Can we construct suitable abstractions such that each language can meaningfully talk about the others without importing an overwhelming amount of detail?*

This seminar brought together a wide variety of researchers working on the above questions. The seminar was organized into

a series of short and long talks on topics related to the above overarching questions, and four breakout sessions focussing on broader questions and challenges. Next, we briefly summarize the talks and sessions. The contributed talks focussed on the following over arching themes – *semantics*, *type systems*, *program analysis*, *contracts*, *languages* and *security*.

4.39 Exploring Interdisciplinary Grand Challenges in ICT Design to Support Proactive Health and Wellbeing

Organizers: m. c. schraefel and Elizabeth F. Churchill

Seminar No. 14272

Date: June 29 to July 2, 2014 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.4.6.108

Participants: Lars L. Andersen, Susanne Boll, Alan Chamberlain, Adrian David Cheok, Elizabeth F. Churchill, Maria Francesca Costabile, Ed Cutrell, Catalina Danis, Adrian Friday, Katherine Isbister, Wolfgang Maaß, Florian Michahelles, Sean Munson, Les Nelson, Erika Poole, Albrecht Schmidt, m. c. schraefel, Katie A. Siek, Gisela Sjogaard, Carlos Trenado

To date poor health costs billions annually, negatively impacting our nations' GDPs. Costs include provision of healthcare for acute and chronic physical and mental conditions and reductions in productivity resulting from absences from work due to sickness.

Much hope has been placed in the deployment of networked information and communications technologies (ICTs) to improve the health of citizens, engage them in proactive healthcare strategies, and thus reduce the likelihood of illness in the first place. Part of the promise is that ICTs in the form of personal, commercial and infrastructural/ governmental platforms may be deployed ubiquitously, pervasively and more cost-effectively than one-to-one human care.

This perspective draws primarily from advances in mHealth and eHealth in the medical community. The focus in these domains, however, is to see health as a medical condition, focussing on tracking and management of patient records, support for doctor-patient interaction, and technologies for regimen adherence and therapy management. In our view, an excellent complement to this perspective is a focus on Proactive Health and Wellbeing, where the concept of health is broadened from being the absence or management of a medical condition or conditions to include a personal engagement with and understanding of wellbeing. ICT has, so far, delivered less success in this arena [1].

The Perspectives Workshop on Exploring Interdisciplinary Grand Challenges in ICT Design to Support Proactive Health and Wellbeing was convened to engage with these issues. We invited scholars to focus on Proactive Health and to elicit what key challenges we need to address in ICT that, if we were to put in concerted and coordinated effort as a community, would have demonstrable effect. We invited reflection on the promise of ICT in contributing to global health, GDP and wellbeing. Our participants have come from various areas in computer science, principally Human Computer Interaction, Data Science and Information Studies, both from industry and academia. We also had participation from psychology, sociology, sports science, medicine and neural science. While most participants were established research leads, we also reached out to up-and-coming, early career researchers in Computer Science and Human Computer Interaction who have a developing track record on health and wellbeing related topics. These individuals will be the future leads in this emerging field.

Over the three days of our workshop we developed 5 key challenge areas, focusing on the significance of each challenge, success at year 1, 3 and 10, as well as resources required to facilitate success. These areas correspond to data sciences, motivational modeling, design thinking, framework building, and a higher order rethinking of the space of "health":

1. Developing Effective Methodologies, Measures and Metrics for Understanding Proactive Health and Wellbeing. Small and "Big" data need to be captured, cleaned and curated to more effectively reflect hard-to-measure experiential aspects of wellbeing. Qualitative data are needed to better understand what is being captured quantitatively, and to enable a deeper understanding of the diversity of experience and to more deeply investigate what is represented in the data within the "long tail".
2. Understanding Motivation and Sensemaking. New models of motivation and sense making are needed in order to more deeply understand people's aspirations and the contingencies of their everyday lives that enable or prevent personal proactive health and wellbeing practices. A move from imposing normative models of "change" to understanding how sustained motivation and self- and other-persuasion can result in new and innovative technology-enabled programs is needed. .
3. Rethinking Design Practices. We need reflective design practices that focus on the phenomenological aspects of a design to complement designs that focus on intervention and

instrumental goal achievement. This arena relates to the need for better motivational models, but addresses the ways in which our design practices mould what we create. How can we more effectively move basic science into applied science and more effective engineering?

4. Creating New Frameworks and Models. We need to develop frameworks and models that take into account unconscious as well as conscious drivers of human behavior, that better connect 'body', 'mind' and 'feeling' experiences, that address emotions as well as cognitive processing, and that acknowledge rhythms of participation and non-participation that are health-positive as well as those that are health negative. This requires a deeper engagement with psychosocial, brain and biological sciences to develop and bring into perspective more holistic frameworks and models.
5. Rethinking the Phenomenology and Epistemology of "Health". Rolling the previous areas up, one of the broader challenges directly addresses how to drive multi disciplinary thinking in regard to proactive wellbeing. A new field of enquiry at the intersection of Human Computer Interaction (HCI) and Computer Science, we need to think about how to motivate and increase engagement from researchers, from designers and engineers, from policy makers, from governmental agencies and from business leaders.

The key outcome of the workshop is an affirmation that a focus on Proactive Health and Wellbeing is both timely and socially necessary, and represents a viable area of research and development. A suite of near-term future activities have been planned and "owned" by participants to drive forward in the coming 6 months. Activities include a follow up Dagstuhl seminar, and workshops, panels, summer schools, invited publications, special issues, and the establishment of an area conference. We have also agreed to explore new ways to engage around experimental design, feedback and collaborative work. We invite potential collaborators to contact us for further discussion and to learn more about our ongoing efforts in this emerging arena of Wellth Sciences.

References

1 William T. Riley, Daniel E. Rivera, Audie A. Atienza, Wendy Nilsen, Susannah M. Allison, and Robin Mermelstein. Health behavior models in the age of mobile interventions: are our theories up to the task? *Translational behavioral medicine*, 1(1):53–71, 2011.

4.40 Feature Interactions: The Next Generation

Organizers: Sven Apel, Joanne M. Atlee, Luciano Baresi, and Pamela Zave

Seminar No. 14281

Date: July 6–11, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.7.1

Participants: Oana M. Andrei, Sven Apel, Joanne M. Atlee, Luciano Baresi, Sandy Beidu, Bruno Cafeo, Marsha Chechik, Gerhard Chroust, Krzysztof Czarnecki, Nicolas Dintzner, Sebastian Erdweg, Kathi Fisler, Stefania Gnesi, Thomas Gschwind, Reiner Hähnle, Michael Jackson, Cliff B. Jones, Christian Kästner, Mario Kolberg, Sergiy Kolesnikov, Shriram Krishnamurthi, Malte Lochau, Oscar M. Nierstrasz, Gilles Perrouin, Christian Prehofer, Gunter Saake, Sandro Schulze, Norbert Siegmund, Stefan Sobernig, Mirco Tribastone, Alexander von Rhein, Andrzej Wasowski, Pamela Zave

Overview and Motivation

A major goal of software and systems engineering is to construct systems from reusable parts, which we call *features* (end-user–visible units of behavior or increments in system functionality). Such a compositional approach can decrease time to market, improve product quality, and diversify the product portfolio. However, the success of a compositional approach depends on the modularity of the reusable parts. The quest for modularity has a long tradition in software and systems engineering, programming languages research, and even in newer fields such as synthetic biology.

In the early days of software and systems engineering, the feature-interaction problem was identified (and coined) as a major threat to modularity [8, 25, 31]. A feature interaction occurs when the behavior of one feature is affected by the presence of another feature. Often the interaction cannot be deduced easily from the intended behaviors of the individual features involved. A canonical example is the inadvertent interaction between the call-forwarding and call-waiting features of a telephony system [8]: If both features are active, the system can reach an undefined possibly unsafe state when it receives a call on a busy line, because it is not specified whether the call should be suspended or forwarded. Alternatively, a feature interaction can be *planned*: for example, advanced cruise-control features are designed to interact with and extend basic cruise control.

To be safe, software developers must analyze the consequences of all possible feature interactions, in order to find and fix the undesired interactions. The *feature-interaction problem* is that the number of potential interactions to consider is exponential in the number of features. As a result, software developers find that their work in developing new features is dominated by the tasks to detect, analyze, and verify interactions.

The feature-interaction problem is deeply rooted in the fact that the world is often not compositional [20, 25]. That is, a feature is not an island. It communicates and cooperates with other features and the environment, so it cannot be completely isolated. Insights from complex-systems research suggest that feature interactions are a form of emergent behavior that is inherent to any system that consists of many, mutually interacting parts. So, emergent system behavior – which is not deducible from the individual parts of a system – can be observed in many situations including in quantum systems (e. g., superconductivity), biological systems (e. g., swarm intelligence), and economical systems (e. g., trading market crashes). The challenge is to foster and manage *desired* interactions and to detect, resolve, and even avoid *undesired* feature interactions – in a scalable manner.

The feature-interaction problem became a crisis in the telecommunications industry in the late 1980s [5]. To handle complexity, there was the strong desire to *compose* systems from independently developed features, but there was no means to detect, express, and reason about feature interactions. Researchers responded with formalisms that enable automatic detection of feature interactions [4, 7, 14, 15, 21, 26], architectures that avoid classes of interactions [17, 18, 28, 29, 31], and techniques for resolving interactions at run-time [16, 27]. Architectural solutions have been the most successful because they impose general coordination strategies (i. e., serial execution) that apply to all features that are 'plugged' into the architecture, thereby, addressing the scalability issue at the heart of the feature-interaction problem. In coordination-based approaches, such as BIP [2, 3] or Composition Patterns [10], the interactions among a set of features are specified explicitly and can be specialized for subsets of features.

While the pioneering work on the feature-interaction problem in telecommunication systems was foundational and very successful [8], it is limited in the sense that it is based on assumptions that hold for telecommunication systems, but that do not hold in other domains. For example, architecture-based approaches take advantage of the fact that communication takes place over a mostly

serial connection between communicating parties – which is not the case in systems made up of parallel components (e. g., service systems, automotive software) or software product lines (e. g., features implemented via conditional compilation such as the Linux kernel). Specifying interactions explicitly is not a general solution either. When facing systems composed of thousands of features, attempting to identify and model a possibly exponential number of feature interactions is elusive. Furthermore, the highly dynamic nature of feature (or service) composition in self-adaptive systems, dynamic product lines, cloud computing, and systems of systems imposes a new class of challenges to solving the feature-interaction problem [1, 9, 24].

So, it is not surprising that different notions of feature interactions have emerged in different communities [6]. Instances of the feature-interaction problem have been observed and addressed in Internet applications [11], service systems [30], automotive systems [12], software product lines [19], requirements engineering [23], computational biology [13], and in many other fields outside of computer science. While all instances of the problem are rooted in the nature of modularity and compositionality [20, 25], the individual views, interpretations, and possible solutions differ considerably. For example, the view on feature interactions taken in program synthesis [22] differs significantly from the view in automotive systems engineering [12]: there are structural vs. behaviour views, static vs. dynamic views, sequential vs. parallel views, functional vs. non-functional, coordinated vs. emergent-behaviour views, and so on. It turns out that feature interactions are a much more general concept than investigated in the past in the context of telecommunication systems, but a classification, comparison, and generalization of the multitude of different views is missing.

The feature-interaction problem is still of pivotal importance in various industrial applications, but, despite significant efforts, it is far from being solved. The underlying hypothesis of organizing a Dagstuhl seminar on this topic was that the time is ripe to gather researchers and practitioners from different areas of computer science and other disciplines to compare, discuss, and consolidate their views, experience, and domain-specific solutions to the feature-interaction problem. To make progress, scientific discourse on the feature-interaction problem must be based on a broader foundation to be able to join forces of different communities. Can other domains learn from the success of domain-specific solutions for telecommunication systems? Are there key principles, patterns, and strategies to represent, identify, manage, and resolve feature interactions that are domain-independent, that are valid and useful across domains? Or, should we strive for domain-specific solutions that are only loosely related to solutions from other domains? Can we develop a unified terminological and conceptual framework for feature-interaction research? Is that even possible or meaningful, given that interactions in telecommunication systems and emergent behavior and phase transitions in swarm systems are, although related, quite different views?

Goals of the Seminar and Further Activities

It is our goal and firm belief that the feature-interaction problem needs to be viewed from a broader perspective. While feature interactions are still a major challenge in software and systems engineering, both in academia and industry, research on the feature-interaction problem has diversified and diverged in the last decade. Researchers working on similar problems, but in different contexts, are largely disconnected and unaware of related work. A major goal of the seminar was to (re)launch a sustained research community that embraces researchers and practitioners from different fields within and outside computer science. We firmly believe that we reached this goal with our seminar. In particular, a subset of the participants is going to organize a follow-up seminar that directly builds on this seminar's results. The next major milestone will be – now as we gained a better understanding of the similarities and differences between the different notions of feature interactions – to establish a catalog on feature-interaction patterns and solutions thereof. The idea for this pattern catalog arose from the final panel session of the seminar. It is inspired by work on patterns in architecture (of buildings). Such a catalog will be the necessary basis for further research on leveraging patterns for detecting, managing, and resolving feature interactions in different kinds of systems.

References

1 L. Baresi, S. Guinea, and L. Pasquale. Service-oriented dynamic software product lines. *IEEE Computer*, 45(10):42–48, 2012.

2 A. Basu, M. Bozga, and J. Sifakis. Modeling heterogeneous real-time components in BIP. In *Proc. of the Int'l Conf. on Software Engineering and Formal Methods (SEFM)*, pages 3–12. IEEE, 2006.

3 S. Bliudze and J. Sifakis. The algebra of connectors – Structuring interaction in BIP. *IEEE Transactions on Computers*, 57(10):1315–1330, 2008.

4 J. Blom, B. Jonsson, and L. Kempe. Using temporal logic for modular specification of telephone services. In *Feature Interactions in Telecommunications Systems*, pages 197–216. IOS Press, 1994.

5 T. Bowen, F. Dworack, C. Chow, N. Griffeth, G. Herman, and Y.-J. Lin. The feature interaction problem in telecommunications systems. In *Proc. of the Int'l Conf. on Software Engineering for Telecommunication Switching Systems (SETSS)*, pages 59–62. IEEE, 1989.

6 G. Bruns. Foundations for features. In *Feature Interactions in Telecommunications and Software Systems VIII*, pages 3–11. IOS Press, 2005.

7 G. Bruns, P. Mataga, and I. Sutherland. Features as service transformers. In *Feature Interactions in Telecommunications Systems V*, pages 85–97. IOS Press, 1998.

8 M. Calder, M. Kolberg, E. Magill, and S. Reiff-Marganiec. Feature interaction: A critical review and considered forecast. *Computer Networks*, 41(1):115–141, 2003.

9 B. Cheng, R de Lemos, H. Giese, P. Inverardi, J. Magee, et al. Software engineering for self-adaptive systems: A research roadmap. In *Software Engineering for Self-Adaptive Systems*, LNCS 5525, pages 1–26. Springer, 2009.

10 S. Clarke and R. Walker. Composition patterns: An approach to designing reusable aspects. In *Proc. of the Int'l Conf. on Software Engineering (ICSE)*, pages 5–14. IEEE, 2001.

11 R. Crespo, M. Carvalho, and L. Logrippo. Distributed resolution of feature interactions for Internet applications. *Computer Networks*, 51(2):382–397, 2007.

12 A. Dominguez. *Detection of Feature Interactions in Automotive Active Safety Features*. PhD thesis, School of Computer Science, University of Waterloo, 2012.

13 R. Donaldson and M. Calder. Modular modelling of signalling pathways and their cross-talk. *Theoretical Computer Science*, 456(0):30–50, 2012.

14 A. Felty and K. Namjoshi. Feature specification and automated conflict detection. *ACM Transactions on Software Engineering and Methodology*, 12(1):3–27, 2003.

15 M. Frappier, A. Mili, and J. Desharnais. Defining and detecting feature interactions. In *Proc. of the IFIP TC 2 WG 2.1 Int'l Workshop on Algorithmic Languages and Calculi*, pages 212–239. Chapman & Hall, Ltd., 1997.

16 N. Griffeth and H. Velthuijsen. The negotiating agents approach to runtime feature interaction resolution. In *Feature Interactions in Telecommunications Systems*, pages 217–235. IOS Press, 1994.

17 J. Hay and J. Atlee. Composing features and resolving interactions. In *Proc. of the ACM SIGSOFT Symp. on Foundations of Software Engineering (FSE)*, pages 110–119. ACM, 2000.

18 M. Jackson and P. Zave. Distributed feature composition: A virtual architecture for telecommunications services. *IEEE Transactions on Software Engineering (TSE)*, 24(10):831–847, 1998.

19 P. Jayaraman, J. Whittle, A. Elkhodary, and H. Gomaa. Model composition in product lines and feature interaction detection using critical pair analysis. In *Proc. of the Int'l Conf. on Model Driven Engineering Languages and Systems (MoDELS)*, LNCS 4735, pages 151–165. Springer, 2007.

20 C. Kästner, S. Apel, and K. Ostermann. The road to feature modularity? In *Proc. of the Int'l Workshop on Feature-Oriented Software Development (FOSD)*, pages 5:1–5:8. ACM, 2011.

21 F. Lin and Y.-J. Lin. A building block approach to detecting and resolving feature interactions. In *Feature Interactions in Telecommunications Systems*, pages 86–119. IOS Press, 1994.

22 J. Liu, D. Batory, and C. Lengauer. Feature oriented refactoring of legacy applications. In *Proc. of the Int'l Conf. on Software Engineering*, pages 112–121. ACM, 2006.

23 A. Nhlabatsi, R. Laney, and B. Nuseibeh. Feature interaction: The security threat from within software systems. *Progress in Informatics*, (5):75–89, 2008.

24 L. Northrop, P. Feiler, R. Gabriel, J. Goodenough, R. Linger, T. Longstaff, R. Kazman, M. Klein, D. Schmidt, K. Sullivan, and K. Wallnau. Ultra-large-scale systems – The software challenge of the future. Technical report, Software Engineering Institute, Carnegie Mellon University, 2006.

25 K. Ostermann, P. Giarrusso, C. Kästner, and T. Rendel. Revisiting information hiding: Reflections on classical and nonclassical modularity. In *Proc. of the Europ. Conf. on Object-Oriented Programming (ECOOP)*, LNCS 6813, pages 155–178, 2011.

26 K. Pomakis and J. Atlee. Reachability analysis of feature interactions: A progress report. In *Proc. of the Int'l Symp. on Software Testing and Analysis (ISSTA)*, pages 216–223. ACM, 1996.

27 S. Tsang and E. Magill. Learning to detect and avoid run-time feature interactions in intelligent networks. *IEEE Transactions on Software Engineering (TSE)*, 24(10):818–830, 1998.

28 G. Utas. A pattern language of feature interaction. In *Feature Interactions in Telecommunications Systems V*, pages 98–114. IOS Press, 1998.

29 R. van der Linden. Using an architecture to help beat feature interaction. In *Feature Interactions in Telecommunications Systems*, pages 24–35. IOS Press, 1994.

30 M. Weiss, B. Esfandiari, and Y. Luo. Towards a classification of web service feature interactions. *Computer Networks*, 51(2):359–381, 2007.

31 Pamela Zave. Modularity in Distributed Feature Composition. In *Software Requirements and Design: The Work of Michael Jackson*, pages 267–290. Good Friends Publishing, 2010.

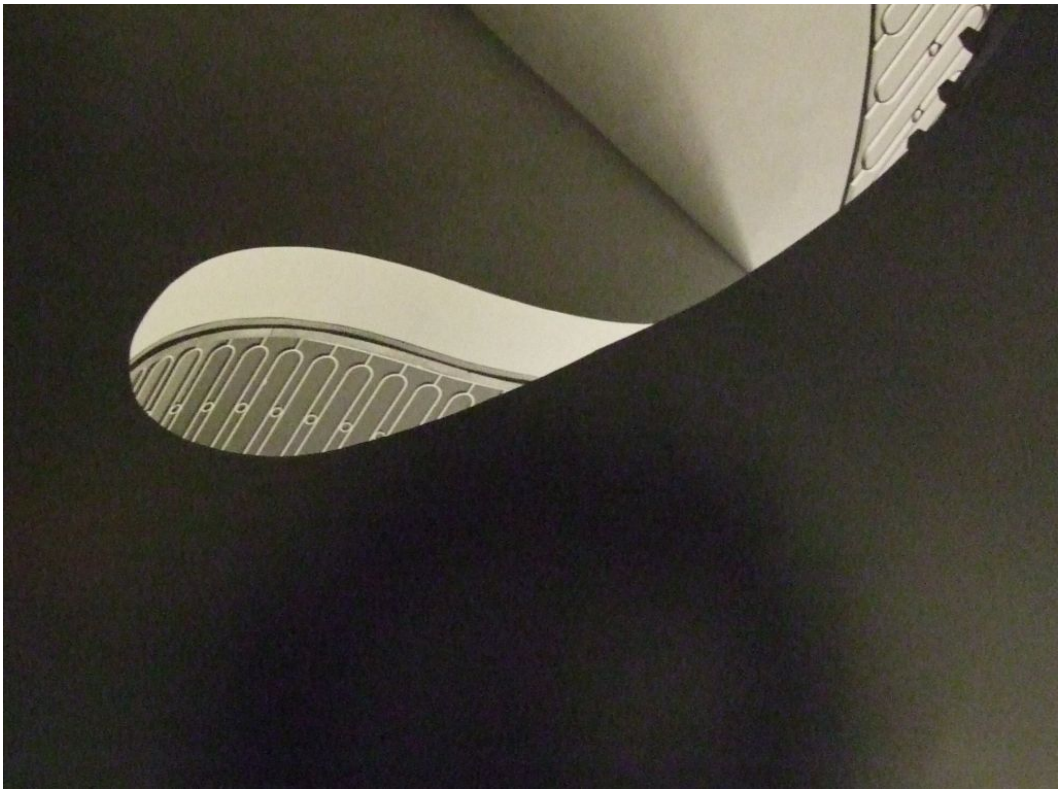

Fig. 4.9
Karola Perrot – Casino 1. Part of the Dagstuhl art collection and donated by: Reinhard Wilhelm and participants in Dagstuhl Seminar 08402.

4.41 Crowdsourcing and the Semantic Web

Organizers: Abraham Bernstein, Jan Marco Leimeister, Natasha Noy, Cristina Sarasua, and Elena Simperl

Seminar No. 14282

Date: July 6–9, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.7.25

Participants: Maribel Acosta, Sofia Angeletou, Lora Aroyo, Abraham Bernstein, Irene Celino, Philippe Cudré-Mauroux, Roberta Cuel, Gianluca Demartini, Michael Feldman, Yolanda Gil, Carole Goble, Robert Kern, Jan Marco Leimeister, Atsuyuki Morishima, Natasha Noy, Valentina Presutti, Marta Sabou, Harald Sack, Cristina Sarasua, Elena Simperl, Markus Strohmaier, Gerd Stumme, Tania Tudorache, Maja Vukovic, Christopher A. Welty, Marco Zamarian

The aim of the *Dagstuhl Seminar 14282: Crowdsourcing and the Semantic Web*, which was held in July 2014, was to gain a better understanding of the dual relationship between crowdsourcing and Semantic Web technologies, map out an emerging research space, and identify the fundamental research challenges that will need to be addressed to ensure the future development of the field.

The seminar focused on three categories of topics: first and foremost we looked into existing crowdsourcing approaches and how these could or have been applied to solve traditional semantic data management tasks. Particular attention was paid to core components of a crowdsourcing-enabled data management and processing system, including methods for quality assurance and spam detection, resources, task and workflow management, as well as interfaces, and the way these components can be assembled into coherent frameworks. A second category of topics that was addressed during the seminar reached out to other disciplines such as economics, social sciences, and design, with the aim to understand how theories and techniques from these fields could be used to build better crowdsourcing-enabled data management systems for the Semantic Web. Last, but not least, we discussed the usage of semantic technologies within generic crowdsourcing scenarios, most notably as means to describe data, resources and specific components.

The seminar, in its community-formative role, represented the starting point for the emergence of working groups that will in the future jointly address the identified scientific challenges. Participants were asked to provide a 1-page position statement reflecting on why they think it makes sense to consider the two topics – crowdsourcing and Semantic Web (or Web of Data) – at the same seminar. Specifically, participants were asked to write a statement reflecting on one of both of the following questions:

1. What are the Semantic Web tasks where you felt you needed crowdsourcing? Why? What were the challenges?
2. What are the crowdsourcing tasks where using semantics might help? Why? What are the challenges?

The first two days of the seminar were dedicated to presentations of topics related to position statements and working groups on use case scenarios and challenges identified during the talks and Q&A sessions. The third day focused on the consolidation of the results of the working groups and the definition of next steps and follow-up activities.

In the following sections we present the position papers written by the researchers of the crowdsourcing and the Semantic Web community, who took part in the seminar. We will publish a more complete research roadmap for crowdsourcing and the Semantic Web at a later stage.

4.42 Information-Centric Networking 3

Organizers: Dirk Kutscher, Taekyoung Kwon, and Ignacio Solis

Seminar No. 14291

Date: July 13–16, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.7.52

Participants: Bengt Ahlgren, Tohru Asami, Kenneth L. Calvert, Antonio Carzaniga, György Dán, Elwyn Davies, Anders Eriksson, Suyong Eum, Kevin R. Fall, Xiaoming Fu, Massimo Gallo, Ashish Gehani, Volker Hilt, Jussi Kangasharju, Holger Karl, Vikas Kawadia, Minyoung Kim, Dirk Kutscher, Taekyoung Kwon, Stefan Lederer, Eiichi Muramoto, Edith Ngai, Börje Ohlman, David Oran, Craig Partridge, Diego Perino, Ioannis Psaras, Damien Saucez, Thomas C. Schmidt, Glenn Scott, Jan Seedorf, Ignacio Solis, Christian Tschudin, Ersin Uzun, Matthias Wählisch, Cedric Westphal, George Xylomenos

Information Centric Networks (ICN) has been a growing area of research in the past few years. The Dagstuhl ICN Seminar series has played a central role in forming the research community. The first seminar, Dagstuhl Seminar 10492, was the meeting point of the various ICN projects across the world; both from the academic perspective as well as the commercial perspective.

The community created at this event continued interacting. It was not long before the members created a set of academic workshops at the most important networking conferences; SIGCOMM, INFOCOM, etc. Following the success of the second Dagstuhl Seminar (12361), the community continued to coalesce and founded the ICNRG. The ICNRG, Information Centric Networking Research Group, was formed at the IRTF to evaluate the technology and to create a forum for companies discuss possible standardization efforts.

With the third iteration of this Seminar we've attempted to bring together the academic and commercial community together once more to discuss the state of the art in ICN. Specifically, we've focused on scalability and deployment. First, what are the problems we face in terms of scaling ICN. Are there technical limitations or political limitations. Second, what are the roadblocks in the path towards deployment. Since there will be no overnight switch, the technology must be deployed in controlled environments where interoperability can be slowly achieved.

4.43 Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures

Organizers: Marc Dacier, Frank Kargl, Hartmut König, and Alfonso Valdes

Seminar No. 14292

Date: July 13–16, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.7.62

Participants: Ali Abbasi, Magnus Almgren, Nils Aschenbruck, Gunnar Björkman, Damiano Bolzoni, Alvaro Cárdenas Mora, Marco Caselli, Jorge R. Cuéllar, Hervé Debar, Sven Dietrich, Ulrich Flegel, Dina Hadziosmanovic, Frank Kargl, Stefan Katzenbeisser, Richard A. Kemmerer, Stephan Kleber, Hartmut König, Marina Krotofil, Pavel Laskov, Michael Meier, Simin Nadjm-Tehrani, Heiko Patzlaff, Andreas Paul, Konrad Rieck, Rene Rietz, Robin Sommer, Radu State, Jens Tölle, Alfonso Valdes, Rens van der Heijden, Alexander von Gernler, Stephen Wolthusen, Emmanuele Zambon

From July 13–16, 2014, more than 30 researchers from the domain of critical infrastructure security met at Schloss Dagstuhl to discuss the current state of security in industrial control systems.

Recent years have highlighted the fact that security precautions of information and communication technology (ICT) in many critical infrastructures are clearly insufficient, especially if considering targeted attacks carried out by resourceful and motivated individuals or organizations. This is especially true for many industrial control systems (ICS) that control vital processes in many areas of industry that are relying to an ever-larger extent on ICT for monitoring and control in a semi or fully automated way. Causing ICT systems in industrial control systems to malfunction can cause huge economic damages or even endanger human lives. The Stuxnet malware that actually damaged around 1000 Uranium enrichment centrifuges in the Iranian enrichment facility in Natanz is the most well-known reported example of an ICT attack impacting ICS.

This situation led to increased efforts in research which also resulted in a number of Dagstuhl seminars related to this topic of which this seminar is a follow-up event, namely two Dagstuhl seminars on "Network Attack Detection and Defense" in 2008 and 2012 and one on "Securing Critical Infrastructures from Targeted Attacks" held in 2012. The main objective of our this latest seminar was to discuss new approaches and ideas on how to detect attacks on industrial control systems and how to limit the impact on the physical components. This is closely coupled to the question of whether and how reactive security mechanisms like Intrusion Detection Systems (IDS) can be made more ICS- and process-aware. To some extent it seems possible to adopt existing security approaches from other areas (e. g., conventional networks, embedded systems, or sensor networks) and one of the questions is whether adopting these approaches is enough to reach the desired security level in the specific domain of industrial control systems, or if approaches specifically tailored for ICS or even single installations provide additional benefits.

The seminar brought together junior and senior experts from both industry and academia, covering different scenarios including electrical grids, but also many other control systems like chemical plants and dike or train control systems. Apart from the detection and prevention of attacks by both security and safety mechanisms, there was an extensive discussion on whether or not such systems should be coupled more strongly from a security perspective. It was also argued that there exists a very diverse space of application domains, many of which have not yet been subject to much study by security researchers, for various reasons. Many of these discussions were triggered by plenary or short talks, covering topics from the state of the art in ICS security, forensics in ICS, security assessments, and the new application domain of flood management.

Apart from talks and subsequent discussions, a number of working groups were organized during the seminar, intended to address specific issues in the field. In total, there were four working groups, each of which provided a summary of their results included in this report. The first was on forensics, discussing how attacks can be detected and analyzed after the fact. A second working group addressed the issue of security and risk management, analyzing why existing IT security approaches do not work for ICS and discussing potential improvements. Industry 4.0 and the wide range of new and non-classical ICS use cases was the topic of a third working group, which discussed the new security challenges arising from these emerging research topics. Finally, there was a working group on the detection of cyber-physical attacks; a core question here were advantages and disadvantages of process-aware intrusion detection mechanisms. The group also discussed the interaction between intrusion detection, intrusion response, and security management.

Based on the talks, discussions and working groups, the

Dagstuhl seminar was closed with a final plenary discussion which summarized again the results from the working groups and led to a compilation of a list of open issues that participants consider necessary to be addressed. Those issues partly overlap with the list of open issues identified in the seminar proposal but also uncovered many new challenges that may become highly relevant research topics and may lead to a new agenda for future research. Those issues are discussed at the end of this report.

4.44 Computational Humanities – Bridging the Gap Between Computer Science and Digital Humanities

Organizers: Chris Biemann, Gregory R. Crane, Christiane D. Fellbaum, and Alexander Mehler

Seminar No. 14301

Date: July 20–25, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.7.80

Participants: Loretta Auvil, David Bamman, Sven Banisch, Chris Biemann, Christopher Brown, Marco Büchler, Gregory R. Crane, Jana Diesner, Christiane D. Fellbaum, Anette Frank, Kurt Gärtner, Siegfried Handschuh, Gerhard Heyer, Fotis Jannidis, Brian Joseph, Cathleen Kantner, Kai-Uwe Kühnberger, Jonas Kuhn, Andy Lücking, Alexander Mehler, David Mimno, Meinard Müller, Andrea Rapp, Geoffry Rockwell, Szymon Rusinkiewicz, Joachim Scharloth, Maximilian Schich, Ute Schmid, Susan Schreibman, David A. Smith, Caroline Sporleder, Wolfgang Stille, Manfred Thaller, Claire Warwick, Katharina A. Zweig

Motivation

Research in the field of *Digital Humanities*, also known as *Humanities Computing*, has seen a steady increase over the past years. Situated at the intersection of computing science and the humanities, present efforts focus on building resources such as corpora of texts, images, musical pieces and other semiotic artifacts digitally available, searchable and analyzable. To this end, computational tools enabling textual search, visual analytics, data mining, statistics and natural language processing are harnessed to support the humanities researcher. The processing of large data sets with appropriate software opens up novel and fruitful approaches to questions in the "traditional" humanities. Thus, the computational paradigm has the potential to transform them. One reason is that this kind of processing opens the way to *new* research questions in the humanities and especially for *different* methodologies for answering them. Further, it allows for analyzing much larger amounts of data in a quantitative and automated fashion – amounts of data that have never been analyzed before in the respective field of research. The question whether such steps ahead in terms of quantification lead also to steps ahead in terms of the quality of research has been at the core of the motivation of the seminar.

Obviously, despite the considerable increase in digital humanities research, a perceived gap between the traditional humanities and computer science still persists. Reasons for this gap are rooted in the current state of both fields: since computer science excels at automating repetitive tasks regarding rather low levels of content processing, it can be difficult for computer scientists to fully appreciate the concerns and research goals of their colleagues in the humanities. For humanities scholars, in turn, it is often hard to imagine what computer technology can and cannot provide, how to interpret automatically generated results, and how to judge the advantages of (even imperfect) automatic processing over manual analyses.

To close this gap, the organizers proposed to boost the rapidly emerging interdisciplinary field of *Computational Humanities* (CH). To this end, they organized a same-named Dagstuhl Seminar that brought together leading researchers in the fields of Digital Humanities and related disciplines. The seminar aimed at solidifying CH as an independent field of research and also at identifying the most promising directions for creating a common understanding of goals and methodologies.

At the core of the organizers' understanding of CH is the idea that CH is a discipline that should provide an algorithmic foundation as a bridge between computer science and the humanities. As a new discipline, CH is explicitly concerned with research questions from the humanities that can more successfully be solved by means of computing. CH is also concerned with pertinent research questions from computing science focusing on multimedia content, uncertainties of digitisation, language use across long time spans and visual presentation of content and form.

In order to meet this *transdisciplinary* conception of CH, it is necessary to rethink the roles of both computer scientist and humanities scholars. In line with such a rethinking, computer scientists cannot be reduced to software engineers whose task is just to support humanities scholars. On the other hand, humanities scholars cannot be compelled to construe post-hoc explanations for results from automatic data analysis. Rather, a common vision – shared among both groups of scientists – is needed that defines and exemplifies accepted methodologies and measures for assessing the validity of research hypotheses in CH. This vision motivated and formed a common ground for all discussions throughout the seminar.

Goals and Content of the Seminar

In order to elaborate the vision of CH as a bridge between computer science and the humanities, the seminar focused on questions that can be subsumed under four different reference points of problematizing CH:

1. **The Present State: What works, what does not?**
   - Review of the success of the last 10 years of the digital humanities: Can we identify commonalities of successful projects? What kinds of results have been obtained? What kinds of results were particularly beneficial for partners in different areas of research? Can success in one field be transferred to other fields by following the same methodology?
   - Review of the challenges of the last 10 years of the digital humanities: What are recurring barriers to efficient cross-disciplinary collaboration? What are the most common unexpected causes of delays in projects? What are common misunderstandings?
   - What is the current role of computer scientists and researchers in the humanities in common projects, and how do these groups envision and define their roles in this interplay?
2. **Computational Challenges in Computational Humanities:**
   - What research questions arise for computational scientists when processing data from the humanities?
   - How can the success of a computer system for humanities data-processing be evaluated to quantify its success?
   - What are the challenges posed by the demands from the humanities? In particular, how can computer scientists convey the notion of uncertainties and processing errors to researchers in the humanities?
3. **Humanities Challenges in Computational Humanities:**
   - What research questions can be appropriately addressed with computational means?
   - How can we falsify hypotheses with data processing support?
   - What is and is not acceptable methodology when one relies on automatic data processing steps?
4. **Common Vision: Algorithmic Foundations of Computational Humanities:**
   - Can we agree on generic statements about the expressivity of the range of algorithms that are operative in the digital humanities and related fields of research?
   - Can we distinguish complexity levels of algorithms in the computational humanities that are distinguished by their conditions of application, by their expressiveness or even explanatory power?
   - Which conditions influence the interpretability of the output generated by these algorithms from the point of view of researchers in the humanities?

The Program

In order to work through our set of goals, the seminar decided for a mixture of talks, working groups and plenary discussions. To this end, four Working Groups (WG) have been established whose results are reported in respective sections of this report:

- The Working Group on *Ethics and Big Data* (members: Bettina Berendt, Chris Biemann, Marco Büchler, Geoffrey Rockwell, Joachim Scharloth, Claire Warwick) discussed a very prominent topic with direct relationships to recent debates about ethical and privacy issues on the one hand and the hype about big data as raised by computer science on the other. One emphasis of the WG was on teaching how to process big data, how this research relates to legal and ethical issues, and how to keep on public dialogs in which such issues can be openly discussed – beyond the narrow focus of the academic community. A central orientation of this discussion was to prevent any delegation of such discussions to closed rounds of experts ("research ethics boards") which do not support open discussions to a degree seen to be indispensable by the WG. The widespread, fruitful and detail-rich discussion of the WG is reported in more detail in the full report of the seminar.
- The Working Group on *Interdisciplinary Collaborations – How can computer scientists and humanists collaborate?* (members: Jana Diesner, Christiane Fellbaum, Anette Frank, Gerhard Heyer, Cathleen Kantner, Jonas Kuhn, Andrea Rapp, Szymon Rusinkiewicz, Susan Schreibman, Caroline Sporleder) dealt with opportunities and pitfalls of cooperations among computer scientists and humanities scholars. The WG elaborated a confusion matrix that contrasts commonplaces and challenges from the point of view of both (families of) disciplines. Ideally, scientists meet at the intersection which challenges both groups of scientists – thereby establishing CH potentially as a new discipline. In any event, this analysis also rules out approaches that reduce either side of this cooperation to the provision of services, whether in terms of computing services or in terms of data provisions.
- The Working Group *Beyond Text* (members: Siegfried Handschuh, Kai-Uwe Kühnberger, Andy Lücking, Maximilian Schich, Ute Schmid, Wolfgang Stille, Manfred Thaller) shed light on approaches that go beyond language in that they primarily deal with non-linguistic information objects as exemplified by artworks or even by everyday gestures. A guiding question of this WG concerned the existence of content-related features of such information objects that can be explored by computational methods. As a matter of fact, corpus building by example of such artifacts is in many cases still out of reach so that computation can hardly access these objects. Seemingly, any success in "computerizing" research methodologies here hinges largely upon human interpretation. Obviously, this is a predestined field of application of human computation with the power of integrating still rather separated disciplines (e. g., musicology, history of art, linguistics etc.).
- The Working Group on *Literature, Lexicon, Diachrony* (members: Loretta Auvil, David Bamman, Christopher Brown, Gregory Crane, Kurt Gärtner, Fotis Jannidis, Brian Joseph, Alexander Mehler, David Mimno, David Smith) dealt with the role of information as stored in large-scale lexicons for any process of automatic text processing with a special focus on historical texts. To this end, the WG started from the role of lexica in preprocessing, the indispensability of accounting for time-related variation in modeling lexical knowledge, the necessity to also include syntactic information, and the field of application of automatic text analysis. Special emphasis was on error detection, correction and propagation. The WG has been concerned, for example, with estimating the impact of lemmatization errors on subsequent procedures such as topic modeling. In support of computational historical linguistics, the WG made several proposals on how to extend lexica (by morphological and syntactical knowledge) and how to link these resources with procedures of automatic text processing.

Part and parcel of the work of these WGs were the plenary sessions in which they had to present their intermediary results

in order to start and foster discussions. To this end, the whole seminar came together – enabling inter-group discussions and possibly motivating the change of group membership. Beyond the working groups, the work of the seminar relied on several plenary talks which partly resulted in separate position papers as published in this report:

- In his talk on *Digital and computational humanities*, Gerhard Heyer shed light on the role of computer science in text analysis thereby stressing the notion of exploring knowledge or text mining. He further showed how these methods give access to completely new research questions in order to distinguish between (more resource-related) *Digital Humanities* and (algorithmic) *Computational Humanities*.
- In his talk, Chris Biemann tackled the field of *Machine Learning* methods from the point of view of their application to humanities data. He clarified the boundedness of these methods in terms of what is called understanding in the humanities. From this point of view, he pleaded for a kind of methodological awareness that allows for applying these methods by clearly reflecting their limitations.
- In their talk on *On Covering the Gap between Computation and Humanities*, Alexander Mehler & Andy Lücking distinguished differences that put apart both disciplines. This includes a methodological, a semiotic and an epistemic gap that together result via an interpretation gap into a data gap. In order overcome these differences, they pleaded for developing what they call hermeneutic technologies.
- In her talk on *Digital Humanities & Digital Scholarly Editions*, Susan Schreibman gave an overview of her work on multimodal, multicodal digital editions that integrate historical, biographical and geographical data. Her talk gave an example of how to pave the way for a people's history in the digital age. To this end, she integrates recent achievements in data mining (most notably network analysis, geospatial modeling, topic modeling and sentiment analysis).
- In his talk on *How can Computer Science and Musicology benefit from each other?*, Meinhard Müller switched the topic of mainly textual artifacts to musical pieces and, thus, to musical artworks. He explained the current possibilities of automatic analysis of musical pieces and demonstrated this by a range of well-known examples of classical music.

This work nicely shows that computational humanities has the goal of covering all kinds of data as currently analyzed and interpreted in the humanities (see also the Working Group *Beyond Text* for such a view).

The seminar additionally included a range of short talks in which participants presented state-of-the-art results of their research: among others, this included talks by Christopher Brown, Anette Frank, Brian Joseph and Szymon Rusinkiewicz. This work nicely provided information about a range of linguistic and multimodal application areas and, therefore, reflected the rich nature and heterogeneity of research objects in the humanities.

A highlight of the seminar was a plenary discussion introduced by two talks given by Gregory Crane and by Manfred Thaller. These talks started and motivated an academic verbal dispute in which, finally, the whole seminar participated in order to outline future challenges of Digital Humanities with impact beyond the border of these disciplines – even onto the society as a whole. Both talks – on *Evolving Computation, New Research Directions and Citizen Science for Ancient Greek and the Humanities* by Gregory Crane and on *The Humanities are about research, first and foremost; their interaction with Computer Science should be too* by Manfred Thaller – opened a broad discussion about the role of humanities among the sciences and their status within the society.

Last, but not least, we should mention two common sessions with a concurrent seminar on Paleography. These sessions, which took place at the beginning and at the end of the seminars, opened an interesting perspective on one particular field that could be counted as a sub-discipline of Computational Humanities. The paleographers met in Dagstuhl for the second time and discussed some of our CH issues previously; it was fruitful to exchange approaches on how to overcome them.

Conclusion

Most of the working groups used their cooperation as a starting point for preparing full papers in which the theme of the group is handled more thoroughly. To this end, the plenary discussed several publication projects including special issues of well-known journals in the field of digital humanities. A further topic concerned follow-up Dagstuhl seminars. The ongoing discussions around the perceived gap between computer science and the humanities and the various proposals from the participants on how to define, bridge or deny this gap made it clear that the seminar addressed a topic that needed discussion and still needs discussion. The talks, panels and working group discussions greatly helped in creating a better mutual understanding and rectifying mutual expectations.

In a nutshell: the participants agreed upon the need to continue the discussion since CH is a young and open discipline.

Fig. 4.10
Drawing for the Dagstuhl children's guest book by Kira (4), daughter of Dagstuhl Seminar 14142 participant Verena Hafner.

4.45 Digital Palaeography: New Machines and Old Texts

Organizers: Tal Hassner, Robert Sablatnig, Dominique Stutzmann, and Ségolène Tarte

Seminar No. 14302

Date: July 20–24, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.7.112

Participants: Orna Almogi, Vincent Christlein, Nachum Dershowitz, Véronique Eglin, Jihad El-Sana, Gernot Fink, Björn Gottfried, Anna Gutgarts-Weinberger, Tal Hassner, Rolf Ingold, Noga Levy, Marcus Liwicki, Josep Lladós, Frederike Neuber, Jean-Marc Ogier, Robert Sablatnig, Joan Andreu Sanchez Peiro, Wendy Scase, Iris Shagrir, Peter A. Stokes, Dominique Stutzmann, Ségolène Tarte, Nicole Vincent, Georg Vogeler

Digital Palaeography emerged as a research community in the late 2000s. Following a successful Dagstuhl Perspectives Workshop on Computation and Palaeography (12382)[16], this seminar focused on the interaction of Palaeography and computerized tools developed in Computer Vision for the analysis of digital images. Given the present techniques developed to enhance damaged documents, optical text recognition or computer-assisted transcription, identification and categorisation of scripts and scribes, the current technical challenge is to develop "new machines", i. e. efficient solutions for palaeographic tasks, and to provide scholars with quantitative evidence towards palaeographical arguments, even beyond the reading of "old texts" (ancient, medieval and early modern documents), which is of interest to the industry, to the wider public, and to the broad community of genealogists.

The identified core issue was to create the conditions of a fluid and seamless communication between Humanities and Computer Sciences scholars in order to advance research in Palaeography, Manuscript Studies and History, on the one hand, and in Computer Vision, Semantic Technologies, Image Processing, and Human Computer Interaction (HCI) systems on the other hand. Indeed, researchers must articulate their respective systems of proof, in order to produce efficient systems that present palaeographical data quickly and easily, and in a way that scholars can understand, evaluate, and trust. To establish fruitful collaborations, it is thus essential to address the "black box" issue, to make a better use of the outreach potential offered by computerized technologies to enrich palaeographical knowledge, and to facilitate the sharing of both the CS and palaeographical methodologies.

This seminar was able to shed light onto two major evolutions between 2012 and 2014; these notable shifts are to do with interdisciplinary communication and with access to "black box" expertise. On the one hand, the notion of "communication" or "bridging the gap" (as expressed by seminar 14301, which took place in conjunction with our own seminar) has become more specific in that issues and problems are now better identified, understood, and expressed. While the two-fold expression "digital palaeography" might lead one to believe that the communication involves only two sorts of actors, it has been expressed in ways clearer than ever that Digital Palaeography as a field is much more complex than a simplistic adjunction of Computer Sciences and Palaeography; indeed CS research, engineering and software development, support and service, linguistics, palaeography, art history, and cultural heritage institutions (Galleries, Libraries, Archives, and Museums – GLAM) all form part of the Digital Palaeography research arena. Good communication requires correct identification of the roles and competence of each actor, and a well-balanced project has to associate/include/foresee the participation of the other actors. It is for example important to clarify that palaeographers are not responsible for copyright or image quality provided by GLAM institutions, in the same way as CS researcher are not responsible for designing interfaces. Within each community, a better understanding of methods and interests of the actors of the other communities is needed to find the right partners (e. g.: keyword spotting is not alignment; writer identification is not script classification). On the other hand, the "black box" issue seems to have been addressed by most teams through the introduction or increase of interactivity of the software tools they presented; interactivity was used not only as a means to produce clear and convincing results, but also to overcome the shortcomings of strictly automatic approaches. In this sense, the reintroduction of "the human into the loop" (or "the use of the users") is part of a process allowing a better understanding on both sides. The "human in the loop" can

[16] http://dx.doi.org/10.4230/DagMan.2.1.14

and should be integrated at all stages, and, even if this need is not always perceived, it is crucial that substantial efforts be dedicated to making implicit assumptions or knowledge explicit. Special attention should be given to avoid the development of tools relying on tautological approaches where tools or datasets incorporate expectations as an underlying (and often implicit) model. In this regard, one cannot overestimate that an unclear result is as important for historians as a clear-cut clustering. In the middle, the "human" gives feedback on preliminary results, enables the enhancement and improvement of the model, as well as creates ground-truth. The display of intermediary results and the integration of user feedback within the process are a welcome solution offered by the latest developments. Likewise, palaeographers have developed new strategies, in their ways of formulating tool requirements or expressing requirements for which they can evaluate the results themselves, regardless of the software being an opaque black-box (P. Stokes, D. Stutzmann, M. Lawo with B. Gottfried).

Overall, this seminar seems to have operated a paradigm shift from black-box issues to trust issues, in the sense that when we first identified black-box issues, we focussed on "computational black boxes", when "human black boxes" are in fact just as problematic. Instead of focussing on computational black-boxes as an issue, we were able to formulate that the important endeavour is that of establishing trust in the respective methodological approaches to the research questions of the research domains. This trust in methodologies is usually mediated by human interactions ("humans in the loop" again!), and the ways in which scholars are able to share an intuitive understanding of their respective expertises with non-experts.

It hence follows that a new (technical) challenge arises, consisting in the creation and implementation of an integrated software tool, web service suite, or environment that would allow users to access and work with extant datasets and tools. The impetus to take up this challenge resides as much in the Humanities as it does in the Computer Sciences. By aggregating the multiple, isolated, specific tools developed by CS researchers through a common access point, digital humanists would support the development of better evaluation metrics and promote a wider use of CS technologies among more traditional Humanities scholars, who could thus become more aware of the existing tools, more autonomous (i. e. less dependant on CS researchers) and thereby empowered. As a reciprocal positive effect, CS researchers could more easily validate their results and gain access to a wider range of annotated datasets. This challenge is also naturally related to trending key concepts such as "interoperability" and "open access". It furthermore engages with the question of the nature of success metrics in the Humanities, where a successful tool is not only the one giving the best results, it is also one enjoying wide acceptance and a large number of users. Improving ergonomics is mandatory, to put the user in the middle and to accumulate a consistent critical mass of annotations (both as feedback and ground-truth).

4.46 Querying and Reasoning Under Expressive Constraints

Organizers: Michael Benedikt, Carsten Lutz, and Balder Ten Cate

Seminar No. 14331

Date: August 10–14, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.1

Participants: Isolde Adler, Antoine Amarilli, Molham Aref, Marcelo Arenas, Michael Benedikt, Meghyn Bienvenu, Manuel Bodirsky, Pierre Bourhis, Cristina Civili, Gaelle Fontaine, Birte Glimm, Tomasz Gogacz, André Hernich, Ian Horrocks, Yazmin Angélica Ibáñez-García, Emanuel Kieronski, Phokion G. Kolaitis, Roman Kontchakov, Tomer Kotek, Markus Krötzsch, Katja Losemann, Carsten Lutz, Florent R. Madelaine, Jerzy Marcinkowski, Wim Martens, Jakub Michaliszyn, Giorgio Orsi, Magdalena Ortiz, Martin Otto, Reinhard Pichler, Andreas Pieris, Lucian Popa, Ian Pratt-Hartmann, Sebastian Rudolph, Vadim Savenkov, Inanc Seylan, Dan Suciu, Balder Ten Cate, Lidia Tendera, Michaël Thomazo, David Toman, Michael Vanden Boom, Grant Weddell, Scott Weinstein, Frank Wolter, Michael Zakharyaschev, Thomas Zeume

Motivation

Query answering in the presence of expressive constraints and logical rules is a topic that has drawn attention from several different research communities. In databases, the interaction of constraints and queries arises in the context of query optimization – for example, how to make use of integrity constraints such as inclusion dependencies and functional dependencies in running a query more efficiently. The topic is also central to the more recent database topics of data integration and data exchange, where constraints are used in the specification of schema mappings. In the area of knowledge representation, the interaction of constraints and queries plays a great role as well – particularly in ontology-based query answering.

The work in these areas is closely related also to another fundamental topic in theoretical computer science, namely decidable fragments of first-order logic. In particular, many of the query answering and query analysis techniques used in recent work within databases and knowledge representation have close links to static analysis of guarded logics, a family of logics that arose out of work by the modal logic and finite model theory communities.

The seminar focused on the convergence of interest of the databases, knowledge representation, and computational logic communities. Its goal was to make visible the connections between these distinct communities, to look at tools and algorithms in one community that can be applied within others, to understand which formalisms and techniques are most promising from the perspective of practical applications, and to propose new ways to combine techniques across communities.

Overview and Outcome

The week started with three overview lectures from well-known authorities in databases, description logics, and decidable fragments of first-order logic. These talks introduced the necessary background for participants and raised research themes that would be explored in later talks. The week then proceeded with a wide-ranging series of talks by participants. In addition to finite model theory, description logics, and databases, there were also talks concerning the interaction of querying problems with constraint satisfaction. The presentations included theoretical work as well as system demonstrations and discussion of practical obstacles to efficient querying with constraints. There were two presentations by participants from industry (IBM and LogicBlox), describing products that implement integrity constraint-based approaches to entity resolution and data analytics, respectively. There was also a presentation on the status of constraint-based reasoning within the W3C endorsed query language SPARQL. In addition to the formal talks, the seminar had an open discussion session, which included a mention of some major open problems and directions to be explored for the communities, as well an attempt at mapping the distinct vocabularies of the different communities.

A main outcome of the discussion was a desire for further interaction between the communities. There were a number of proposals put forward for how to achieve this, including co-location of a KR-related conference with a database conference like VLDB or SIGMOD/PODS. Another outcome was a collection of topics that were particularly worth pursuing by all communities. The handling of inconsistency in databases was one of these – both further investigation of the most widely-used approach for inconsistency-handling, based on repair and consistent query answering, and the examination of alternative approaches. The notion of repair tied into the question of investigating the relationship of data uncertainty and constraints. Markov logic networks (MLNs) are likely to play a role in reconciling "hard" integrity constraints with probabilities, although the interplay of probabilistic data and classical approaches to integrity constraints will involve a more general revision of the major computational

problems with uncertainty in mind. Another topic identified for future work was the notion of incremental checking of constraints. Incremental computation was alluded to in several talks, but there appears to be a need to take a more holistic look at models for incremental computation and their application in constraint maintenance. The recent activity within dynamic complexity makes the topic of incremental computation within constraint handling particularly ripe for revisiting.

Conclusion

We believe that the seminar was very successful in bringing together the involved communities and in promoting interaction and exchange between them. Similarities as well as differences between the communities' research efforts became clearly visible and the participants conceived the seminar as a significant step forwards in bridging the gap and raising mutual awareness. Many participants expressed interest in a followup event.

4.47 Formal Methods for Coordinating Multi-Agent Systems

Organizers: Thomas Ågotnes and Nils Bulling

Seminar No. 14332

Date: August 10–14, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.21

Participants: Thomas Ågotnes, Markus Brill, Jan M. Broersen, Nils Bulling, Natalia Criado, Mehdi Dastani, Louise A. Dennis, Catalin Dima, Valentin Goranko, Maaike Harbers, Andreas Herzig, Max Knobbout, Brian Lam, Beishui Liao, Brian Logan, Marin Lujak, Nir Oren, Julian Padget, Truls Pedersen, Joshua Sack, Marija Slavkovik, Bas J. G. Testerink, Nicolas Troquard, Paolo Turrini, Birna Van Riemsdijk, Marc Van Zee, Laurent Vercouter

Formal methods form an active and broad field of research in multi-agent systems, ranging from bottom-up to top-down approaches. Properties of individual agents, e. g., aspects related to decision making and knowledge representation, are rather low-level, while the specification and verification of multi-agent systems are higher-level. In particular, *logic-based approaches* have been successfully used for the modeling of intelligent agents and for reasoning about them: epistemic logics allow to talk about knowledge; temporal logics to reason about the evolution of actions; and strategic logics have been proposed to reason about abilities of agents and coalitions. Alternating-time temporal logics and STIT logics are prominent members of the latter type, and more expressive logics like Strategy Logic have recently been proposed. What they all have in common is their descriptive flavor. Typically, they are not used to actively change the state of the agent system but to talk and to reason about the system. Multi-agent logics are particularly relevant for the coordination problem. The latter is concerned with global properties of a system. Since the global behaviour of a system emerges from the individual behaviour of agents, it is not obvious what the global properties are. By specifying global properties using multi-agent logics, verification techniques can be employed to verify what of these properties are met by the system; thus, to find out what the global properties are. Interaction between rational decision makers in general, and coordination problems in particular, have been studied in game theory for decades. However, game theory is not concerned with *computational* or *logical* aspects of coordination: how we can represent and reason about coordination in computers. In contrast, many *agreement technologies* are used in an interactive way, e. g., for arriving at agreements about joint actions or coalition structures. Techniques like norms and social laws coordinate the agent's behavior and often require less interaction of agents with their peers. Agents have to decide whether to comply with the rules or not. A difficult problem is the synthesis of appropriate norms and social laws. Related issues important for appropriate control techniques include the detection of norm violations and sanctioning mechanisms.

The seminar aimed at opening up new directions of research into the coordination problem, by bringing together researchers working in different areas of multi-agent systems as well as related fields, and in particular, to combine insights from research in the following fields:

- formal methods and verification, and multi-agent logics in particular,
- game theory in multi-agent systems, and
- agreement technologies.

The seminar took place between 10 and 14 August, 2014. This medium-size, four day seminar was highly international: the 27 participants came from 12 different countries. The scientific program consisted of presentations, discussions and working groups. We scheduled presentations of three different types: overview, medium, and short. The aim of the four one hour overview talks was to give a broad introduction of the main fields relevant to the seminar – to provide a common ground. They covered Argumentation Theory, Normative Systems, Judgement Aggregation, and Computational Social Choice. Then, we scheduled ten medium (20 minutes long) and ten short (15 minutes long) presentations. We encouraged the speakers to give rather informal, non conference-style talks focussing on high-level ideas in order to provide input for the discussion groups.

From the discussions, two working groups emerged which focused on one of the following topics:

- Concepts: conceptual definition and classification – what is coordination, coordination problems, and solutions?
- Formalisation of coordination

We organized three meetings for the working groups and two

joint discussion sessions for presenting and discussing the results of the working groups.

In addition to the scientific program, we enjoyed a hike which was followed a Barbecue, and the unique atmosphere of Dagstuhl.

4.48 Resource-bounded Problem Solving

Organizers: Yll Haxhimusa, Iris van Rooij, Sashank Varma, and H. Todd Wareham

Seminar No. 14341

Date: August 17–22, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.45

Participants: Burcu Arslan, Tarek R. Besold, Mark Blokpoel, Sarah Carruthers, Christopher Cherniak, Nicole Cruz de Echeverria Loebell, Harmen de Weerd, Michael R. Fellows, Nadine Fleischhut, Liane Gabora, Emmanuel Genot, Nina Gierasimczuk, Vinod Goel, Yll Haxhimusa, Justine Jacot, Frank Jäkel, Brendan Juba, Alexandra Kirsch, Etienne Koechlin, Antonina Kolokolova, Johan H. P. Kwisthout, Falk Lieder, Wolfgang Maass, Matthias Mnich, Martin Möhrmann, David Noelle, Stellan Ohlsson, Maria Otworowska, Zygmunt Pizlo, Frances A. Rosamond, Constantin Rothkopf, Zahra Sajedinia, Ulrike Stege, Marieke Sweers, Niels A. Taatgen, John K. Tsotsos, Iris van Rooij, Sashank Varma, Rineke Verbrugge, Todd Wareham, Scott Watson

This Dagstuhl Seminar on 'Resource-bounded Problem Solving' was a successor to Dagstuhl Seminar 11351: 'Computer Science & Problem Solving: New Foundations', held in August 2011, which was the first Dagstuhl event to bring together computer scientists and psychologists to exchange perspectives on problem solving in general. Before summarizing the content of the seminar itself, we describe the theoretical motivations for the topic of 'Resource-bounded Problem Solving', and the choice for the interdisciplinary composition of participants, ranging complexity theory, cognitive psychology, artificial intelligence and cognitive neuroscience.

Background and Motivation

Problem solving, whether by humans or machines, is bounded by the resources at hand. For machines, these resources fundamentally include hardware and processing speed. For humans, important resources also include mental representations, memory capacities, inferential capacities, and time to name a few. All these resources can be severely limited, by constraints imposed by both the implementing hardware (current technology in the case of machines, the organization of our brains in the case of humans) and the physical and social operating environments.

The study of resource-bounded problem solving has a long and productive history within computer science, which has resulted in a number of efficient exact-solution algorithms and algorithm design techniques as well as, where such algorithms are not possible, various widely-applicable approximate-solution heuristics. Given that brains have evolved to solve problems under severe resource constraints, resource-bounded problem solving may also provide one of the best windows on the organization of cognitive brain function. Trying to model exactly how humans solve really hard—that is, resource demanding—problems efficiently seems a good strategy from the perspective of deriving predictive and explanatory cognitive models. After all, many cognitive models may be able to match human performance on really easy problems, but only a few can for really hard problems. Hence, if one finds even one cognitive model that can solve a hard problem as well as humans one can be much more confident that it captures fundamental principles of human cognition.

What makes tasks resource demanding? Here computational complexity theory is a key source of information for the study of human and machine problem solving. Computational complexity theory studies the intrinsic resource demands of various computational problems. It allows us to assess when resource demands are low, reasonable, high, or impractically high. Though the importance of computational complexity has been recognized in computer science for decades, it has been underutilized to date by cognitive scientists. This is not for want of opportunities, as cases are known where cognitive scientists have studied principles of resource-bounded problem solving in apparent ignorance of relevant computational complexity results. The following example illustrates such a situation.

- Solving constrained Traveling Salesman problems: MacGregor and Ormerod (1996,Attention, Perception & Psychophysics) hypothesized that the difficulty of solving Euclidean versions of the Traveling Salesperson problem (E-TSP) may be more a function of the number of inner points (i. e., the points not lying on the convex hull of the point set) than of the total number of points. This hypothesis was tested and confirmed for human subjects solving pen-and-paper instances of E-TSP. Independently, it was shown that E-TSP is fixed-parameter tractable when the number of inner points is the parameter (Deineko et al.; COCOON 2004). In other words, it is possible to solve E-TSP in time $f(k)n^c$, where fs is a non-polynomial function of the number k of inner points, n is the total number of points, and c is a constant. This result is relevant for explaining the findings of MacGregor and

Ormerod (1996) as it gives a computational formalization of their hypothesis.

There may be many other such opportunities waiting to be noticed. There is also evidence that when such opportunities are exploited and cognitive scientists and complexity theorists establish collaborations, these collaborations can yield novel perspectives and approaches. Below are two examples of such ongoing collaborations and their products.

- Analyzing resource demands of cognitive models: Using computational complexity concepts and techniques, psychologists can systematically study how human problem solving proceeds under various resource demands. Also, complexity theory can predict how resource demands scale as a function of a problem's parameters. Psychologists can then in turn use these predictions to test the models being studied. This approach has been successfully implemented by Iris van Rooij, Todd Wareham, and others in a wide variety of domains, including decision-making (2005, Journal of Mathematical Psychology) and analogical problem solving (2011, Journal of Problem Solving). This program has also led to the development of a theory of structure approximability which has produced new results within both computer science (2007, Proceedings of Dagstuhl seminar 07281) and psychology (2012, Journal of Mathematical Psychology).
- Pyramid data structures and efficient search: Humans and animals are able to delineate, detect and recognize objects in complex scenes at a blink of an eye. Tsotsos (1990, Behavioral and Brain Sciences) performed a complexity analysis and showed that hierarchical representation of visual information and hierarchical processing of this information is one of the best, if not the best, way for brains to solve visual problems. Pyramid data structures provide an effective model for efficient hierarchical search of the problem space. This perspective has led to fruitful collaborations between Yll Haxhimusa, Zygmunt Pizlo, Walter Kropatsch and others, yielding new algorithmic techniques in computer vision (2005, Pattern Recognition Letters; 2009, Vision & Computing),as well as inspiring new cognitive models of visual problem solving in psychology (2006 and 2011, Journal of Problem Solving).

With this seminar we aimed to actively stimulate the exchange of ideas and results between computational complexity theorists, psychologists, cognitive neuroscientists, and AI-researchers studying problem solving. In particular, we wanted to ensure that this exchange would be of use to each (and not just one) of these communities. We believe that such n-way productivity is crucial to fruitful long-term interdisciplinary collaboration, in that it encourages the continued interest of members of all communities in collaborating.

Organization of seminar

On Day 1 of the seminar, Iris van Rooij opened the seminar by explaining its history, motivation and aims. This was following by a round of introductions, in which each participant introduced themselves, who they are, what their home disciplines were, what their relevant research interests were, and what they hoped to both bring to the seminar and get out off it.

Given the interdisciplinary nature of the questions of interest and the wide range of expertise of the seminar participants, it was crucial that a common understanding of the different goals and assumptions of the disciplines involved at the meeting be established early in the meeting. To this end, the first full day of the seminar was devoted to primers on basic terminology, assumptions, and goals of four major sub disciplines represented at the seminar (namely, computational complexity theory, artificial intelligence, psychology, and cognitive neuroscience).

The four introductory keynote speakers, Zygmunt Pizlo (Cognitive Psychology), Todd Wareham (Computational Complexity Theory), Rineke Verbrugge (Artificial Intelligence), and David Noelle (Cognitive Neuroscience), all addressed each of the following questions for their own respective fields.

- What are the goals of that discipline?
- What are the techniques used in that discipline?
- What do the terms "problem solving" and "resource bounds" mean in that discipline?
- What does that discipline have to offer to other disciplines in the context of this seminar?
- What does that discipline want from the other disciplines in the context of this seminar?
- What are some tentative research questions and collaboration opportunities?

Following these introductory key notes there was a panel discussion. Day 1 closed with a Town Hall meeting, in which the set-up and organization for the next days was discussed with all participants and a preliminary schedule was established (later on, as needed, this schedule was updated). At the Town Hall meeting the concept of Birds-of-a-Feather (B.O.F.) sessions was also explained, which turned out to be a very successful format for self-organized working groups.

Days 2, 3 and 4 involved a mix of talks, posters, B.O.F. sessions, and Town Hall sessions. Pairs or triplets of talks were followed by panel discussions to stimulate cross talk connections. Poster sessions allowed for more in-depth discussion in an informal setting, and B.O.F. sessions allowed people to gather and discuss more specific topics of common interest. At Town Hall sessions, a plenary report on the insights gained from each B.O.F. session was given, so that all participants were kept up to date of the outcome of such events. All in all, this set-up worked very well, stimulating active exchange and discussion between participants that crossed disciplinary boundaries.

On the morning of Day 5 we closed with reflections on the overall organization and content of the seminar. It was a shared perspective among participants that the seminar had been exceptionally successful in bringing together the different fields involved in the seminar, initiating many first-time theoretical exchanges and conceptualizations of possible common research questions. Many participants indicated that following this seminar, they would be interested in more focused seminars specializing in subtopics within the domain of problem solving or specializing in specific modeling techniques.

4.49 Equilibrium Computation

Organizers: Nimrod Megiddo, Kurt Mehlhorn, Rahul Savani, and Vijay V. Vazirani

Seminar No. 14342

Date: August 17–22, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.73

Participants: Ilan Adler, Susanne Albers, Xi Chen, Richard Cole, Rachel Cummings, Kousha Etessami, Martin Gairing, Jugal Garg, Paul W. Goldberg, Tobias Harks, Sergiu Hart, Penny Haxell, Martin Hoefer, Kamal Jain, Nimrod Megiddo, Kurt Mehlhorn, Ruta Mehta, Peter Bro Miltersen, Dimitris Paparas, Britta Peis, Yuval Rabani, Rahul Savani, Guido Schäfer, Leonard J. Schulman, Alexander Skopalik, Madhu Sudan, László A. Végh, Bernhard von Stengel

The aim of this seminar was to study research issues related to algorithms and complexity for computation of equilibria in games and markets. The majority of participants were academics from computer science departments; some were from other disciplines; and several participants were from the corporate research departments of eBay, IBM, and Microsoft. All participants have strong interdisciplinary interests that typically span Economics, Game Theory, and Theoretical Computer Science.

The seminar started with a session of lightening talks, in which participants had two minutes and one slide to introduce themselves. This session was extremely well received, and it was worth the effort to ensure that everyone submitted a slide in advance. It is an effective and efficient way for everyone to get to know a little bit about each other, and thus to have things to talk about outside of talks right from the start of the seminar.

Three tutorials were given on topics chosen by the organizers. Bernhard von Stengel gave a tutorial on complementary pivoting algorithms for the Linear Complementarity Problem (LCP). The tutorial focussed on geometric aspects of LCPs and complementary pivoting algorithms, and in particular Lemke's algorithm. The LCP captures many game and market problems, and it came up throughout the seminar, most directly in the final talk by Adler on reductions to bimatrix games from PPAD Lemke-verified LCPs.

Complementary pivoting algorithms inspired the complexity class PPAD, which, together with FIXP, capture the problems of finding fixed points and equilibria of games and markets. The second tutorial, given by Kousha Etessami, was about the complexity of equilibria and fixed points. It covered PPAD (= linear-FIXP), FIXP, and FIXP-a, and discussed some associated open problems. Related contributed talks included the following. Etessami, in a separate talk, showed that the complexity of computing a (perfect) equilibrium for an n-player extensive form game of perfect recall is hard for FIXP-a. Gairing showed that the problem of finding an equilibrium of a weighted congestion game is FIXP-hard. Garg presented several results on market equilibria, including the result that it is FIXP-hard to compute an equilibrium of an Arrow-Debreu exchange market with Leontief utility functions. Chen presented a PPAD-hardness result for the problem of finding an approximate equilibrium in an anonymous game with seven actions per player. Mehta showed that it is PPAD-hard to find an equilibrium of a rank-3 bimatrix game. Paparas presented PPAD-hardness results for several market settings with non-monotone utilities. The number of talks related to these complexity classes shows their ongoing importance for the field of equilibrium computation.

The third tutorial was on game dynamics and was given by Sergiu Hart. He showed that "uncoupledness" severely limits the possibilities to converge to Nash equilibria, but on the other hand, there are simple adaptive heuristics, such as "regret matching", that lead to correlated equilibria. At the end of his tutorial, Hart also presented an exponential lower bound on the query complexity of correlated equilibria. In a closely related contributed talk, Goldberg gave bounds for the query complexity of approximate equilibria of various types, including for the relatively new concept of ϵ-well-supported correlated equilibrium.

A large number of contributed talks presented algorithms for computing equilibria of games and markets. On market equilibria we had the following algorithmic talks: Cole presented an asynchronous gradient descent method that implements asynchronous tâtonnement; Mehlhorn presented a combinatorial polynomial-time algorithm for the linear exchange model; Vazirani introduced Leontief-Free Utility Functions and presented a complementary pivoting algorithm for computing an equilibrium in markets with these utilities; and Vegh presented new convex programmes for linear Arrow-Debreu markets. On other game models, we had the following algorithmic talks: Cummings presented an efficient

differentially private algorithm for computing an equilibrium in aggregative games; Savani presented a gradient descent algorithm for finding an approximate equilibrium of a polymatrix game; and Skopalik presented algorithms for finding approximate pure equilibria of congestion games.

There were other contributed talks on a range of topics: Harks talked about resource competition on integral polymatroids; Hoefer talked about decentralized secretary algorithms; Jain presented an analysis of several business models and pricing schemes; and Schäfer presented results about coordination games on graphs.

Apart from the topics of the tutorials, all other talk topics were chosen by the presenters, not by the organizers. Generally talks were informal, and were very interactive, often with lengthy discussions taking place during them. All talks were well received. Open problems were discussed in two sessions, the first during a normal seminar room session, and the second with cheese and wine in the evening. Below we give abstracts for the talks and brief summaries of the open problems that were discussed.

4.50 Decision Procedures and Abstract Interpretation

Organizers: Daniel Kroening, Thomas W. Reps, Sanjit A. Seshia, and Aditya V. Thakur

Seminar No. 14351

Date: August 24–29, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.89

Participants: Aws Albarghouthi, Joshua Berdine, Régis Blanc, Martin Brain, Jörg Brauer, Alessandro Cimatti, Vijay D'Silva, Bruno Dutertre, Alberto Griggio, Arie Gurfinkel, Julien Henry, Jacob Howe, Daniel Kroening, Akash Lal, Antoine Miné, Ruzica Piskac, Thomas W. Reps, Michel Rueher, Peter Schrammel, Sanjit A. Seshia, Rahul Sharma, Rohit Sinha, Aditya Thakur, Charlotte Truchet, Thomas Wies, Hongseok Yang, Florian Zuleger

The seminar was successful in bringing the following two communities together:

- designers and implementors of abstract interpreters, and
- designers and implementors of decision procedures.

The abstract interpretation (AI) and decsion procedure (DP) communities have several interests in common. Tools created by each of these communities can be viewed as using symbolic techniques to explore the state state of a transition system. However, the respective repertoires of techniques used in the two disciplines are quite different, and each community has its own mindset and outlook. The seminar sought to capitalize on recent ideas that demonstrated new connections between the two disciplines, and, consequently, promote the cross-fertilization between the areas at a deep technical level.

The seminar had 27 participants from both the AI and DP communities. To keep pariticipants from both areas engaged during a session, the organizers refrained from filling a session only with talks focusing on a particular community. Instead, each session consisted of talks by participants of both research communities. Furthermore, talks by young researchers were scheduled earlier in the week, which enabled them to get better feedback on their research.

The Dagstuhl Seminar 14352 "Next Generation Static Software Analysis Tools" was held concurrently with Dagstuhl Seminar 14351. There were a number of joint activities organized to foster interaction among participants of the two seminars:

- The first session on Monday was a joint session for participants of both seminars. In this session, all participants introduced themselves and briefly described their research interests. Furthermore, Patrick Cousot, an organizer for Seminar 14352, and Thomas Reps, an organizer for Seminar 14351, each gave a "scene-setting" talk.
- The Wednesday excursion to the steel mill and Egyptian exhibit was organized as a joint activity.
- A joint session was organized on Thursday afternoon. The talks in this session were given by participants of both seminars.
- The seating arrangement for the Friday dinner was organized so that participants from both seminars sat together.
- The schedule of talks for both seminars was shared with all participants. Hence, participants of one seminar were able to attend a specific talk in the other seminar, if they felt the talk was especially relevant.

Apart from the planned activities listed above, the week saw a lot of informal discussions among participants of these two seminars in the evenings.

The seminar also featured talks about two other research areas: constraint programming (CP) and machine learning (ML). The talks by Mine, Rueher, and Truchet highlighted the use of abstract interpretation in CP. The talks by Reps, Seshia, Sharma, and Thakur discussed the application of ML techniques, such as inductive learning, to problems in AI and DP. Both these sets of talks garnered interesting discussions about the connections among all these various reseach areas. Furthermore, this discussion indicates that future seminars should include even more researchers and practitioners from not just the AI and DP communities, but also the CP and ML communities.

Fig. 4.11
Petra Stoeppel – Durch die Stadt. Part of the Dagstuhl art collection and donated by Saarland Sporttoto GmbH.

4.51 Next Generation Static Software Analysis Tools

Organizers: Patrick Cousot, Daniel Kroening, and Carsten Sinz

Seminar No. 14352

Date: August 24–29, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.107

Participants: Roberto Bagnara, Dirk Beyer, Mehdi Bouaziz, Patrick Cousot, Tomasz Dudziak, David Faragó, Christian Ferdinand, Jérôme Feret, Marcelo Frias, Vijay Ganesh, Roberto Giacobazzi, Christoph Gladisch, Udo Gleich, Manuel Hermenegildo, Ralf Huuck, Daniel Kroening, K. Rustan M. Leino, Francesco Logozzo, Peter Müller, Filip Niksic, Andreas Podelski, Hendrik Post, Francesco Ranzato, Xavier Rival, Helmut Seidl, Carsten Sinz, Michael Tautschnig, Shmuel Tyszberowicz, Caterina Urban, Tomas Vojnar, Milena Vujosevic-Janicic, Reinhard Wilhelm

Software errors are still a widespread plague. They manifest themselves, e. g., in program crashes, malfunction, incorrect behavior, or security vulnerabilities. Even software that has been in use for decades and has been deployed to millions of users (e. g., the compression library zlib) still contains flaws that are revealed only now and have to be fixed. Both in academia and industry considerable effort has been undertaken to develop tools and methodologies to obtain fault-free software. Nowadays, static analysis tools, which search for program errors without running the software, have reached a state where they are, in some industries (e. g., the automotive or avionics industry), already part of the standard software development and quality assurance process (with tools and companies like, e. g., Polyspace, Coverity, KlocWork, AbsInt, or Astrée). And although these tools can help finding residual errors more quickly, they still suffer from some shortcomings:

- Lack in precision. For a certain fraction of program locations in the source code it cannot be decided whether there is an error or not. Such "undecided cases" require (often time-consuming) manual rework, limiting the value of such tools.
- Due to the manual effort required, static software analysis tools have not yet made their way to mainstream software development (besides industries, where software reliability is indispensable and considerable amounts of time and money are spent on quality assurance).

Over the last years, software analysis tools based on abstract interpretation have been refined and tools based on new core formalisms, such as model checking, have gained traction, mainly in the form of two key methods: counterexample-guided abstraction refinement (CEGAR), and bounded model checking (BMC). The success of these new tools was, to a substantial part, enabled by the enormous progress that was made on the underlying logical decision procedures (SAT and SMT solvers). New software analysis tools based on these techniques come with considerably improved precision (less false positives), but they are still not competitive with tools based on abstract interpretation with respect to scalability. Also, they are rarely used in industrial software development projects so far.

With this seminar we believe that we were able to stimulate further progress in this field by intensifying the collaboration between (a) researchers on new static software analysis tools, (b) scientists working on improved high-performance decision procedures, and (c) practitioners, who know what is needed in industry and which kind of software analysis tools are accepted by developers and which are not.

The Dagstuhl Seminar was attended by participants from both industry and academia. It included presentations on a wide range of topics such as:

- Recent trends in static analysis, consisting of new algorithms and implementation techniques.
- New decision procedures for software analysis, for example, to analyze programs with complex data structures.
- Industrial case studies: What are the problems industrial users of static analysis tools are facing?
- Experience reports and statements on current challenges.

The first day of the seminar started with an introduction round, in which each participant shortly presented his research interests. As the seminar was held concurrently with a second, closely related Dagstuhl Seminar on "Decision Procedures and Abstract Interpretation" (14351), the introductory session was held jointly by both seminars. Four overview talks were also organized jointly by both seminars, and were given by Thomas Reps, Patrick Cousot, Vijay Ganesh, and Francesco Logozzo.

There was also a tool demonstration session on Thursday afternoon, in which seven tools were presented (15 minutes each).

In further talks of the seminar young as well as senior researchers presented on-going and completed work. Tool developers and participants from industry reflected on current challenges in the realm of software analysis.

The seminar was concluded with a panel discussion about the current challenges of static software analysis for industrial application.

We expect that with this Dagstuhl Seminar we were able to make a step forward towards bringing static software analysis tools to every programmer's workbench, and therefore, ultimately, improve software quality in general.

4.52 Computational Aspects of Fabrication

Organizers: Marc Alexa, Bernd Bickel, Sara McMains, and Holly E. Rushmeier

Seminar No. 14361

Date: August 31 to September 5, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.126

Participants: Marc Alexa, Moritz Bächer, Bernd Bickel, Nathan Carr, Duygu Ceylan, Yong Chen, Paolo Cignoni, Stelian Coros, Bailin Deng, Piotr Didyk, Gershon Elber, Wolfgang Heidrich, Kristian Hildebrand, Leif Kobbelt, Manfred Lau, Sylvain Lefebvre, Steve Marschner, Sara McMains, Eder Miguel, Niloy Mitra, Stefanie Müller, Mark Pauly, Nico Pietroni, Helmut Pottmann, Holly E. Rushmeier, Szymon Rusinkiewicz, Ryan Schmidt, Carlo H. Sequin, Kristina Shea, Karan Singh, Melina Skouras, Bernhard Thomaszewski, Nobuyuki Umetani, Gordon Wetzstein, Tim Weyrich, Emily Whiting, Denis Zorin

As manufacturing goes digital, the current understanding of industrial production will change fundamentally[17]. The digital age in manufacturing is coupled with new output devices that allow rapid customization and rapid manufacturing, revolutionizing the way we design, develop, distribute, fabricate, and consume products. We need to find computational models that support this new way of production thinking and lead its technological understanding. This opens challenges for many areas of science research, such as material science, chemistry, and engineering, but also and perhaps foremost computer sciences. The currently available digital content creation pipelines, algorithms, and tools cannot fully explore new manufacturing capabilities. To meet these demands, we need a deep understanding of computer graphics fundamentals: Shape, appearance of shape and materials, and physically-based simulation and animation. When designing an object, there is an inherent interplay among all these fundamental aspects.

The purpose of this seminar is to bring together leading experts from academia and industry in the area of computer graphics, geometry processing, and digital fabrication. The goal is to address fundamental questions and issues related to computational aspects of fabrication and jump-start collaborations that will pioneer new approaches in this area.

[17] Special report: manufacturing and innovation. *The Economist* 403(8781):46, 2012.

Fig. 4.12
Drawing for the Dagstuhl children's guest book by Isidora (15 months), daughter of Dagstuhl Seminar 14301 participant David Smith and partner Cynthia Mentzer.

4.53 Augmenting Human Memory – Capture and Recall in the Era of Lifelogging

Organizers: Mark Billinghurst, Nigel Davies, Marc Langheinrich, and Albrecht Schmidt

Seminar No. 14362

Date: August 31 to September 5, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.8.151

Participants: C. Philip Beaman, Michel Beaudouin-Lafon, Agon Bexheti, Mark Billinghurst, Andreas Bulling, Ozan Cakmakci, Nigel Davies, Simon Dennis, Tilman Dingler, Christos Efstratiou, Niforatos Evangelos, Scott Greenwald, Cathal Gurrin, Vicki Hanson, Niels Henze, James D. Hollan, Christoph Korinke, Kai Kunze, Wendy E. Mackay, Mateusz Mikusz, Daniela Petrelli, Adam C. Rule, Corina Sas, Albrecht Schmidt, Aurélien Tabard, Kristof Van Laerhoven, Geoff Ward, Katrin Wolf

Technology has always had a direct impact on how and what humans remember. This impact is both inevitable and fundamental – technology radically changes the nature and scale of the cues that we can preserve outside our own memory in order to trigger recall. Such change is not new – we have seen the transition from story-telling to written books, from paintings to photographs to digital images and from individual diaries to collective social networks. However, in recent years three separate strands of technology have developed to the extent that collectively they open up entirely new ways of augmenting human memory:

1. near-continuous collection of memory cues has become possible through the use of technologies such as Microsoft's SenseCam, social networks and interaction logs.
2. advances in data storage and processing now enables widespread mining of stored cues for proactive presentation, both in terms of cues collected by an individual and in terms of complex networks of related cues contributed by others.
3. the presence of ubiquitous displays (both in the environment and via personal devices such as Google Glasses) provides many new opportunities for displaying memory cues to trigger recall.

It is self-evident that we do not effectively encode all of the information that we encounter, nor are we able to retrieve at will, all of that content that we do encode. When trying to recall known facts, many of our day-to-day memory failures result from a temporary failure to retrieve memories rather than from their permanent erasure. Our ability to recall target information is particularly vulnerable to transient changes in accessibility that arise through changes in the contents of our short-term memories and the cues in our environment. That memory can be improved with effective cues is beyond doubt: whilst a typical participant might be able to recall only 38 out of a set of 100 words that had been continually studied and sorted over many minutes, this accuracy increases to 96% when the most effective cues are presented at test. One experiences these temporary failures to retrieve memories everyday when we might remark "I cannot recall X (e. g., his name, the company, the town, etc.), but if I saw it, I would recognise it". Tellingly, we are unlikely to experience or say the converse.

One of the most frustrating features of human memory is that we are particularly vulnerable at remembering to do something in the future (the area of memory research known as prospective memory). Prospective memory failures readily occur for remembering time-based future events (hence the value of setting computer alarms reminding us of meetings), and for remembering event-based future events (remember to post a letter on the way to work, remember to pick up a takeaway for the family tonight). Research suggests that whereas there is a general decline in memory with increasing old age, it is prospective memory and retrieval in the absence of cues that are particularly impaired, whereas cued recall and recognition are more preserved.

The Dagstuhl Seminar 14362 "Augmenting Human Memory – Capture and Recall in the Era of Lifelogging" focused on a vision of the world in which augmented memory systems make everyday use of peripheral, ambient multi-media content – delivered via large wall-mounted displays, smartphone wallpapers, or wearable in-eye projectors – to intelligently integrate, display, and enable the review of life-relevant personal data. Such memory augmentation technologies have the potential to revolutionise the way we use memory in a wide range of application domains.

Behaviour Change: Effecting behaviour change is an important objective in many important areas such as health (e. g. lifestyle changes such as increasing exercise or stopping smoking) and sustainable transport (e. g. encouraging people to make more environmentally-friendly transport choices). Unfortunately, despite good intentions, many people experi-

ence difficulty in implementing planned behaviour: for example, it is well known that many people are reluctant to make a trip to the gym despite paying large gym membership fees. Psychological theory stresses that intentional behaviours are more likely to be implemented when individuals are reminded of their own attitude towards such behaviours (e. g., the positive gains that will result), and the attitudes of significant others to the behaviour (what loved ones, family, friends, peers, and society in general think of the behaviour and its outcomes). In addition, realistic scheduling is important: planned behaviour is more likely to be performed if it is timetabled with the transition from immediately preceding activities in mind. Finally, behaviour is more likely if it is perceived to be more achievable and more enjoyable. Memory augmentation can help with the realistic scheduling and reminding of the planned activities, and can remind people at the point at which decision making is necessary (e. g., at the planned time to visit the gym) of the positive benefits from the behaviour, the previous good experience of the behaviour and the progress that is being made.

Learning: Such technologies can be used as part of a learning environment. In particular, through the use of ambient displays it might be possible to cue recall, and hence reinforce learning of a wide range of skills. For example, the acquisition of a new language could be supported by providing appropriate cues to facilitate recall of vocabulary. Similarly, a class teacher could be encouraged to remember the names of their pupils, and a study abroad student could learn culturally-significant facts as they explore a new city.

Supporting Failing Memories: Research has shown that as we age, our ability to perform uncued recall is particularly vulnerable to age-related decline. Memory augmentation technologies could be used to help remedy this memory loss by providing older users with time-relevant and context-appropriate cues. In this way, older individuals could enjoy greater self-confidence and greater independence by being reminded of moment-by-moment situated details of where they were, what they were intending to do, and how they could get home. They may also enjoy better relationships if they could be reminded of the autobiographical details of their loved ones (such as the names and ages of their loved ones' children), or if they could review and then be reminded of the details of a recent conversation or event (e. g., a recent day out or family gathering).

Selective Recall: Through appropriate selection of memory cues that are presented to the user, memory augmentation technologies might also be used to facilitate selective recall. According to the psychological theory of retrieval-induced forgetting, the act of reviewing memories not only enhances the probability of spontaneously retrieving these reviewed memories in the future, but it can also attenuate the spontaneous retrieval of related but unreviewed memories. The study of retrieval-induced forgetting has largely been confined to the laboratory using lists of categorised words. It is of both pure and applied interest (e. g., the desired attenuation of unwanted, outdated, or traumatic memories; and the undesired attenuation of wanted but unreviewed memories) to see if this phenomenon can be observed when reviewing a subset of "real world" memories, and if so, we will be able to measure the extent to which unreviewed memories could be attenuated through selective reviewing.

Memory Based Advertising: While many of the application domains for memory augmentation technologies are for the public good, the same technologies can also be employed in the commercial context. For example, such technologies could be used to support a new form of advertising in which users have memories triggered explicitly to drive purchasing decisions. For example, when passing a shop selling luggage a cue could be presented that causes a passer-by to remember a specific experience from their past in which their own luggage didn't work satisfactorily. This may then cause the user to enter the shop and purchase some new luggage.

Collectively, the seminar participants explored the scientific foundations for a new technology eco-system that can transform the way humans remember in order to measurably and significantly improve functional capabilities while maintaining individual control. At its heart lies the creation of memory augmentation technology that provides the user with the experience of an extended and enhanced memory. Such technology is based on recent improvements in the collection, mining, and presentation of appropriate information to facilitate cued memory recall. This research is inherently multidisciplinary and combines elements of pervasive computing, information retrieval and data privacy with psychology and sociology.

4.54 Adjoint Methods in Computational Science, Engineering, and Finance

Organizers: Nicolas R. Gauger, Michael Giles, Max Gunzburger, and Uwe Naumann

Seminar No. 14371

Date: September 8–12, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.9.1

Participants: Christian Bischof, Luise Blank, David Bommes, Martin Bücker, Bruce Christianson, Mike Dewar, Boris Diskin, Patrick Farrell, Christian Frey, Nicolas R. Gauger, Andreas Griewank, Stefanie Günther, Ralf Hannemann-Tamas, Laurent Hascoet, Patrick Heimbach, Paul D. Hovland, Barbara Kaltenbacher, Peter Korn, Drosos Kourounis, Ben Lenser, Johannes Lotz, Marcus Meyer, Viktor Mosenkis, Jens-Dominik Mueller, Uwe Naumann, Eric Nielsen, Jacques Peter, Olivier Pironneau, Alex Pothen, Thomas Rung, Ekkehard Sachs, Max Sagebaum, Stephan Schmidt, Volker Schulz, Markus Towara, Stefan Ulbrich, Jörn Ungermann, Jean Utke, Andrea Walther, Qiqi Wang

The human desire for meaningful numerical simulation of physical, chemical, biological, economical, financial (etc.) phenomena in CSEF has been increasing with the growing performance of the continuously improving computer systems. As a result of this development we are (and will always be) faced with a large (and growing) number of highly complex numerical simulation codes that run at the limit of the available HPC resources. These codes often result from the discretization of systems of PDE. Their run time correlates with the spatial and temporal resolution which often needs to be very high in order to capture the real behavior of the underlying system. There is no doubt that the available hardware will always be used to the extreme. Improvements in the run time of the simulations need to be sought through research in numerical algorithms and their efficient implementation on HPC architectures.

Problem sizes are often in the billions of unknowns; and with emerging large-scale computing systems, this size is expected to increase by a factor of thousand over the next five years. Moreover, simulations are increasingly used in design optimization and parameter identification which is even more complex and requires the highest possible computational performance and fundamental enabling algorithmic technology. Derivatives of certain objectives of these numerical models with respect to a potentially very large number of model parameters are crucial for the highly desirable transition from pure simulation to optimization. Approximation of these derivatives via finite difference quotients often lacks the required accuracy. More importantly, it may be infeasible for a large parameter space in terms of its computational complexity. Adjoint numerical programs have until recently been written by hand to overcome this problem. Such programs compute (large) gradients with machine accuracy at a small constant multiple of the computational complexity of the underlying primal simulation. Due to the enormous size of most numerical simulation codes the manual procedure may take up to several man years. Moreover manual adjoint codes are error-prone and hard to maintain as the primal simulation evolves. Computer scientists have been developing special software tools based on the principles of algorithmic differentiation (AD) to generate discrete adjoint code automatically. Consequently, this method has gained considerable acceptance within the CSEF community as illustrated by numerous successful case studies presented in the proceedings of so far six international conferences on AD. See http://www.autodiff.org for details.

Illustrative Example: Classical applications of adjoint methods arise in the context of large-scale inverse problems, such as the estimation of unknown or uncertain parameters of implementations of mathematical models for real-world problems as computer programs. Imagine the optimization of the shape of an aircraft with the objective to maximize its lift. The continuous mathematical domain (the surface of the aircraft) is typically discretized through the generation of a mesh with a potentially very large number of points spread over the whole surface. Optimization aims to adapt the position of these points in 3D space such that the objective is met while at the same time satisfying various constraints (e. g. prescribed volume). A naive approach might run a potentially very large number of primal numerical simulations with changing mesh configurations thus being able to identify an optimum within this very limited search space.

Derivative-based approaches use information on the sensitivity of the objective at the given mesh configuration with respect to changes in the positions of all mesh points (the gradient) in order to make a deterministic decision about the next configuration to be considered. The sensitivities can be approximated through local perturbations of the position of each mesh point (finite difference quotients). A single optimization step would thus require a number of primal simulations that is of the order of the number of degrees of freedom (three spatial coordinates for each mesh point)

induced by the mesh. This approach is practically infeasible as a single simulation may easily run for several minutes (if not hours) on the latest HPC architectures. The approximation of a single gradient would take months (if not years) for a mesh with only one million points.

Adjoint methods deliver the gradient at the cost of only a few (between 2 and 10) primal simulations. Continuous adjoint methods derive an adjoint version of the primal mathematical model analytically followed by the numerical solution of the resulting adjoint model. While this approach promises low computational cost (approx. 2 primal simulations) it can be mathematically challenging and numerically inconsistent when compared with the primal numerical simulation. To the best of our knowledge, the automation of the derivation of continuous adjoint models is still outstanding.

Discrete adjoint methods rely on the algorithmic differentiation of the primal numerical model, thus overcoming the potential numerical inconsistencies induced by the continuous adjoint. Depending on the mode of implementation of AD, the level of maturity of the AD tool, and the expertise of the user of the tool the computational cost can range between 2 and 20 primal simulations, sometimes even more. Still this cost is independent of the number of mesh points (referring to the above example). Solutions to problems arising in adjoint methods require expertise in both theoretical and applied Computer Science as well as in Numerical Analysis. Robust methods for the data flow reversal within adjoint code are built on special graph partitioning and coloring algorithms. Their implementation on modern HPC architectures (e. g. using MPI and/or OpenMP) has impact on the simulation software design and the data management. The use of accelerators has been considered only recently with many open as of yet unsolved problems. Static and dynamic program analysis and compiler construction techniques have been developed to facilitate the semi-automatic generation of discrete adjoint code. The exploration of a potential extension of these techniques to continuous adjoint code was one of the subjects of this seminar. Other conceptual problems discussed included functional analytic aspects of adjoint methods and their impact on practical implementation, combinatorial problems in adjoint code generation and their computational complexities, and simulation software engineering guidelines in the light of adjoint methods.

Adjoint methods borrow from a variety of subfields of Computer Science and Applied Mathematics including high performance and combinatorial scientific computing, program analysis and compiler construction, functional analysis, numerical analysis and linear algebra, and with relevance to a wide range of potential areas of application. As such, the topic lends itself to a series of seminars taking more detailed looks into the respective subjects. With this seminar we intent to initiate a sequence of related events alternating in between the Leibniz Center for Informatics at Schloss Dagstuhl and the Mathematisches Forschungsinstitut Oberwolfach, thus, emphasizing the obvious synergies between Computer Science and Mathematics in the given context.

4.55 Analysis of Algorithms Beyond the Worst Case

Organizers: Maria-Florina Balcan, Bodo Manthey, Heiko Röglin, and Tim Roughgarden

Seminar No. 14372

Date: September 7–12, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.9.30

Participants: Spyros Angelopoulos, Sanjeev Arora, Maria-Florina Balcan, Shai Ben-David, Markus Bläser, Avrim Blum, Karl Bringmann, Kamiel Cornelissen, Radu Curticapean, Sanjoy Dasgupta, Matthias Englert, Michael Etscheid, Uriel Feige, Tobias Friedrich, Anupam Gupta, Marvin Künnemann, Tengyu Ma, Konstantin Makarychev, Yury Makarychev, Bodo Manthey, Matthias Mnich, Heiko Röglin, Tim Roughgarden, Or Sheffet, Christian Sohler, Ruth Urner, Sergei Vassilvitskii

The theory of algorithms has traditionally focused on worst-case analysis. This focus has led to both a deep theory and many beautiful and useful algorithms. There are, however, a number of important problems and algorithms for which worst-case analysis does not provide useful or empirically accurate results. For example, worst-case analysis suggests that the simplex method is an exponential-time algorithm for linear programming, while in fact it runs in near-linear time on almost all inputs of interest. Worst-case analysis ranks all deterministic caching algorithms equally, while in almost all applications some algorithms (like least-recently-used) are consistently superior to others (like first-in-first-out).

The problem is that worst-case analysis does not take into consideration that worst-case inputs are often rather contrived and occur hardly ever in practical applications. It led to the situation that for many problems the classical theory is not able to classify algorithms meaningfully according to their performance. Even worse, for some important problems it recommends algorithms that perform badly in practice over algorithms that work well in practice only because the artificial worst-case performance of the latter ones is bad.

Only in recent years a paradigm shift towards a more realistic and robust algorithmic theory has been initiated. The development of a more realistic theory hinges on finding models that measure the performance of an algorithm not only by its worst-case behavior but rather by its behavior on typical inputs. However, for an optimization problem at hand it is usually impossible to rigorously define the notion of "typical input" because what such an input looks like depends on the concrete application and on other indistinct parameters. The key to building a rigorous and realistic theory is hence not to define exactly what a typical input looks like, but to identify common properties that are shared by real-world inputs. As soon as such properties are identified, in many cases one can explain why certain heuristics work well in practice while others do not. The next step is then to look for algorithmic means to exploit these properties explicitly in order to obtain improved algorithms in practice that are not tailored to unrealistic worst-case inputs.

Many different models that go beyond classical worst-case theory have been suggested. These models can be divided into two main categories: either they are based on the assumption that inputs are to some extend random (probabilistic analysis) or they consider only inputs that satisfy certain deterministic properties that mitigate the worst case.

Probabilistic Analysis

Average-case analysis is probably the first thought that springs to mind when mentioning probabilistic input models. In such an analysis, one considers the expected performance on random inputs. Starting in the seventies, many algorithms that showed a remarkable performance in practice have been analyzed successfully on random inputs. This includes algorithms for classical optimization problems, such as the traveling salesman problem, or the simplex method for linear programming.

While average-case analysis has been successfully applied to many problems, a major concern is that inputs chosen completely at random have for many problems little in common with inputs arising in practice. Similar as a random TV screen produced by static noise has nothing to do with a typical TV screen, a set of random points does not resemble, for instance, a realistic clustering instance with mostly well-separated clusters.

Smoothed Analysis. To overcome the drawbacks of average-case and worst-case analysis, the notion of smoothed analysis has been suggested by Spielman and Teng in 2001. In this model, inputs are generated in two steps: first, an adversary chooses an arbitrary instance, and then this instance is slightly perturbed at random. The smoothed performance of an

algorithm is defined to be the worst expected performance the adversary can achieve. This model can be viewed as a less pessimistic worst-case analysis, in which the randomness rules out pathological worst-case instances that are rarely observed in practice but dominate the worst-case analysis. If the smoothed running time of an algorithm is low (i. e., the algorithm is efficient in expectation on any perturbed instance) and inputs are subject to a small amount of random noise, then it is unlikely to encounter an instance on which the algorithm performs poorly. In practice, random noise can stem from measurement errors, numerical imprecision or rounding errors. It can also model arbitrary influences, which we cannot quantify exactly, but for which there is also no reason to believe that they are adversarial.

The framework of smoothed analysis was originally invented to explain the practical success of the simplex method for linear programming. Spielman and Teng analyzed linear programs in which each coefficient is perturbed by Gaussian noise with standard deviation σ. They showed that the smoothed running time of the simplex method is bounded polynomially in the input size and $1/\sigma$. Hence, even if the amount of randomness is small, the expected running time of the simplex method is polynomially bounded. After its invention smoothed analysis has attracted a great deal of attention and it has been applied in a variety of different contexts, e. g., in multi-objective optimization, local search, clustering, and online algorithms. By now smoothed analysis is widely accepted as a realistic alternative to worst-case analysis.

Semi-random Models. Semi-random input models can be considered as analogues of smoothed analysis for graph problems and they even predate smoothed analysis by a couple of years. There is a variety of semi-random graph models that go beyond the classical Erdős-Rényi random graphs. In most of these models graphs are generated by a noisy adversary – an adversary whose decisions (whether or not to insert a particular edge) have some small probability of being reversed. Another well-studied class of semi-random models are *planted models*, in which a solution (e. g., an independent set or a partitioning of the vertices in color classes) is chosen and then edges are added randomly or by an adversary of limited power in such a way that the given solution stays a valid solution for the given problem.

Similar to smoothed analysis, semi-random models have been invented in order to better understand the complexity of NP-hard graph problems because Erdős-Rényi random graphs often do not reflect the instances one encounters in practice – many graph problems are quite easy to solve on such random graphs.

Deterministic Input Models

Smoothed analysis and semi-random models are multi-purpose frameworks that do not require much information about how exactly typical inputs for the optimization problem at hand look like. If more information is available, it makes sense to identify structural properties of typical inputs that make them easier to solve than general inputs. There are well known examples of this approach like the TSP, which gets easier (in terms of approximation) when restricted to inputs in which the distances satisfy the triangle inequality. Also in computational geometry it is a very common phenomenon that problems become easier if one assumes that no angles are too small or not too many objects overlap in the same region.

In recent years there has been an increased interest in more sophisticated deterministic input models, in particular for clustering problems. Balcan, Blum, and Gupta introduce and exploit the so-called $(1+\alpha, \varepsilon)$-approximation-stability property of data in the context of clustering. This assumption is motivated by the observation that in many clustering applications there is usually one correct but unknown target clustering, and the goal is to find a clustering that is close to this target clustering and misclassifies only a few objects. On the other hand in the common mathematical formulation of clustering problems a potential function is defined that assigns a value to each clustering. Then a clustering is computed that approximately optimizes the potential function (exact optimization is usually NP-hard). This approach makes sense only if clusterings that approximately optimize the potential function are close to the target clustering. Hence, an implicit assumption underlying this approach is that every clustering that approximately optimizes the objective function is close to the desired target clustering. Balcan et al. made this assumption explicit: they define that a clustering instance satisfies the $(1+\alpha, \varepsilon)$-approximation-stability assumption if in every c-approximation of the potential function at most an ε-fraction of all objects is misclassified compared to the target clustering. Balcan et al. showed that clustering instances with this property are easier to solve than general instances. They have shown specifically how to get ε-close to the target even for values of α for which finding a $1+\alpha$ approximation is NP-hard. Voevodoski et al. have shown that this approach leads to very efficient and accurate algorithms (with improved performance over previous state-of-the-art algorithms) for clustering biological datasets.

Bilu and Linial and later Awasthi, Blum, and Sheffet have considered instances of clustering problems that are *perturbation resilient* in the sense that small perturbations of the metric space do not change the optimal solution. They argue that interesting instances of clustering problems are stable and they prove that the assumption of stability renders clustering polynomial-time solvable. Balcan and Liang further relaxed this assumption to require only that the optimal solution after the perturbations is close to the optimal solution for the unperturbed instance.

These results have triggered a significant amount of work in the past years in the context of clustering and machine learning problems more generally, including subsequent works that proposed new related stability conditions (e. g., the "proximity condition" by Kannan and Kumar). Such works are very good examples demonstrating that identifying properties of real-world inputs can be extremely beneficial for our understanding of algorithmic problems.

Program of the Seminar

The program of the seminar consisted of 23 talks, including the following survey talks:

- Preprocessing of NP-hard problems, Uriel Feige;
- Approximation-stability and Perturbation-stability, Avrim Blum;
- Computational Feasibility of Clustering under Clusterability Assumptions, Shai Ben-David;
- Parametrizing the easiness of machine learning problems, Sanjoy Dasgupta;
- Linear Algebra++: Adventures and Unsupervised Learning, Sanjeev Arora.

The rest of the talks were 30-minute presentations on recent research of the participants. The time between lunch and the afternoon coffee break was mostly left open for individual discussions and collaborations in small groups. One open-problem session was organized.

One of the main goals of the seminar was to foster collaborations among the researchers working in the different branches of analysis of algorithms as sketched above. This is particularly important because at the moment the two communities dealing with probabilistic analysis and deterministic input models are largely disjoint. The feedback provided by the participants shows 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.

Fig. 4.13
Maja Andrack – Photograph on canvas. Detail from invitation to the Maja Andrack art exhibit held at Schloss Dagstuhl on October 14 – December 18, 2014.

4.56 Neural-Symbolic Learning and Reasoning

Organizers: Artur S. d'Avila Garcez, Marco Gori, Pascal Hitzler, and Luís C. Lamb

Seminar No. 14381

Date: September 14–19, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.9.50

Participants: Tsvi Achler, Jim Benvenuto, Tarek R. Besold, Srikanth Cherla, Claudia d'Amato, Artur d'Avila Garcez, James Christopher Davidson, Leo de Penning, Luc De Raedt, Natalia Díaz-Rodríguez, Dominik Endres, Jacqueline Fairley, Jerry A. Feldman, Peter Földiak, Manoel Franca, Christophe D. M. Gueret, Biao Han, Pascal Hitzler, Steffen Hölldobler, Thomas Icard, Randal A. Koene, Kai-Uwe Kühnberger, Luis Lamb, Francesca Alessandra Lisi, Dragos Margineantu, Vivien Mast, Risto Miikkulainen, Andrey Mokhov, Bernd Neumann, Günther Palm, Alan Perotti, Gadi Pinkas, Subramanian Ramamoorthy, Luciano Serafini, Daniel L. Silver, Son Tran, Joshua Welch, Mark Wernsdorfer, Thomas Wischgoll

Neural-symbolic computation aims at building rich computational models and systems through the integration of connectionist learning and sound symbolic reasoning [1, 2]. Over the last three decades, neural networks were shown effective in the implementation of robust large-scale experimental learning applications. Logic-based, symbolic knowledge representation and reasoning have always been at the core of Artificial Intelligence (AI) research. More recently, the use of deep learning algorithms have led to notably efficient applications, with performance comparable to those of humans, in particular in computer image and vision understanding and natural language processing tasks [3–5]. Further, advances in fMRI allow scientists to grasp a better understanding of neural functions, leading to realistic neural-computational models. Therefore, the gathering of researchers from several communities seems fitting at this stage of the research in neural computation and machine learning, cognitive science, applied logic, and visual information processing. The seminar was an appropriate meeting for the discussion of relevant issues concerning the development of rich intelligent systems and models, which can, for instance integrate learning and reasoning or learning and vision. In addition to foundational methods, algorithms and methodologies for neural-symbolic integration, the seminar also showcase a number of applications of neural-symbolic computation.

The meeting also marked the 10th anniversary of the workshop series on neural-symbolic learning and reasoning (NeSy)[18], held yearly since 2005 at IJCAI, AAAI or ECAI. The NeSy workshop typically took a day only at these major conferences, and it became then clear that given that the AI, cognitive science, machine learning, and applied logic communities share many common goals and aspirations it was necessary to provide an appropriately longer meeting, spanning over a week. The desire of many at NeSy to go deeper into the understanding of the main positions and issues, and to collaborate in a truly multidisciplinary way, using several applications (e. g. natural language processing, ontology reasoning, computer image and vision understanding, multimodal learning, knowledge representation and reasoning) towards achieving specific objectives, has prompted us to put together this Dagstuhl seminar marking the 10th anniversary of the workshop.

Further, neural-symbolic computation brings together an integrated methodological perspective, as it draws from both neuroscience and cognitive systems. In summary, neural-symbolic computation is a promising approach, both from a methodological and computational perspective to answer positively to the need for effective knowledge representation, reasoning and learning systems. The representational generality of neural-symbolic integration (the ability to represent, learn and reason about several symbolic systems) and its learning robustness provides interesting opportunities leading to adequate forms of knowledge representation, be they purely symbolic, or hybrid combinations involving probabilistic or numerical representations.

The seminar tackled diverse applications, in computer vision and image understanding, natural language processing, semantic web and big data. Novel approaches needed to tackle such problems, such as lifelong machine learning [6], connectionist applied logics [1, 2], deep learning [4], relational learning [7] and cognitive computation techniques have also been extensively analyzed during the seminar. The abstracts, discussions and open problems listed below briefly summarize a week of intense scientific debate, which illustrate the profitable atmosphere provided by the Dagstuhl scenery. Finally, a forthcoming article describing relevant challenges and open problems will be published at the Symposium on Knowledge Representation and Reasoning:

[18] After the seminar, the *Neural-Symbolic Learning and Reasoning Association* (NeSy) was established; see http://www.neural-symbolic.org/.

Integrating Symbolic and Neural Approaches at the AAAI Spring Symposium Series, to be held at Stanford in March 2015 [8]. This article also adds relevant content and a view of the area, illustrating its richness which may indeed lead to rich cognitive models integrating learning and reasoning effectively, as foreseen by Valiant [9].

Finally, we see neural-symbolic computation as a research area which reaches out to distinct communities: computer science, neuroscience, and cognitive science. By seeking to achieve the fusion of competing views it can benefit from interdisciplinary results. This contributes to novel ideas and collaboration, opening interesting research avenues which involve knowledge representation and reasoning, hybrid combinations of probabilistic and symbolic representations, and several topics in machine learning which can lead to both the construction of sound intelligent systems and to the understanding and modelling of cognitive and brain processes.

References

1 Artur S. d'Avila Garcez, Luis C. Lamb, and Dov M. Gabbay, Neural-Symbolic Cognitive Reasoning. Cognitive Technologies, Springer, 2009.

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

3 D. C. Ciresan, U. Meier, J. Schmidhuber. Multi-column Deep Neural Networks for Image Classification. IEEE Conf. on Computer Vision and Pattern Recognition CVPR 2012.

4 G. E. Hinton, S. Osindero, and Y. Teh, A fast learning algorithm for deep belief nets. Neural Computation, 18, 1527–1554, 2006..

5 Abdel-rahman Mohamed, George Dahl & Geoffrey Hinton. Acoustic Modeling Using Deep Belief Networks. IEEE Transactions on Audio, Speech, and Language Processing. 20(1):14–22, 2012.

6 D. Silver, Q. Yang, and L. Li, Lifelong machine learning systems: Beyond learning algorithms. Proceedings of the AAAI Spring Symposium on Lifelong Machine Learning, Stanford University, AAAI, March, 2013, pp. 49–55.

7 Stephen Muggleton, Luc De Raedt, David Poole, Ivan Bratko, Peter A. Flach, Katsumi Inoue, Ashwin Srinivasan: ILP turns 20 – Biography and future challenges. Machine Learning, 86(1):3–23, 2012.

8 Artur d'Avila Garcez, Tarek R. Besold, Luc de Raedt, Peter Foeldiak, Pascal Hitzler, Thomas Icard, Kai-Uwe Kuehnberger, Luis C. Lamb, Risto Miikkulainen, Daniel L. Silver. Neural-Symbolic Learning and Reasoning: Contributions and Challenges. Proceedings of the AAAI Spring Symposium on Knowledge Representation and Reasoning: Integrating Symbolic and Neural Approaches, Stanford, March 2015.

9 L. G. Valiant, Knowledge Infusion: In Pursuit of Robustness in Artificial Intelligence. FSTTCS, pp. 415–422, 2008.

4.57 Algebra in Computational Complexity

Organizers: Manindra Agrawal, Valentine Kabanets, Thomas Thierauf, and Christopher Umans

Seminar No. 14391

Date: September 21–26, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.9.85

Participants: Farid Ablayev, Manindra Agrawal, Eric Allender, Vikraman Arvind, Markus Bläser, Andrej Bogdanov, Harry Buhrman, Sourav Chakraborty, Arkadev Chattopadhyay, Stephen A. Fenner, Michael Forbes, Lance Fortnow, Anna Gál, William Gasarch, Frederic Green, Rohit Gurjar, Venkatesan Guruswami, Valentine Kabanets, Marek Karpinski, Neeraj Kayal, Pascal Koiran, Swastik Kopparty, Arpita Korwar, Michal Koucký, Matthias Krause, Klaus-Jörn Lange, Sophie Laplante, Meena Mahajan, Or Meir, Peter Bro Miltersen, Natacha Portier, Atri Rudra, Chandan Saha, Rahul Santhanam, Ramprasad Saptharishi, Uwe Schöning, Ronen Shaltiel, Amir Shpilka, Florian Speelman, Amnon Ta-Shma, Thomas Thierauf, Jacobo Torán, Christopher Umans, Nikolay K. Vereshchagin, Ryan Williams, Amir Yehudayoff, David Zuckerman

The seminar brought together almost 50 researchers covering a wide spectrum of complexity theory. The focus on algebraic methods showed the great importance of such techniques for theoretical computer science. We had 25 talks, most of them lasting about 40 minutes, leaving ample room for discussions. In the following we describe the major topics of discussion in more detail.

Circuit Complexity

This is an area of fundamental importance to Complexity. Circuit Complexity was one of the main topics in the seminar. Still it remains a big challenge to prove strong upper and lower bounds. However, the speakers reported amazing progress in various directions.

Or Meir talked on one of the major open problems in complexity theory: proving super-logarithmic lower bounds on the depth of circuits. That is, separating the log-depth circuit class NC^1 from polynomial time, P. Karchmer, Raz, and Wigderson suggested an approach to this problem. The *KRW-conjecture* states that the circuit depth of two functions f and g adds up when we consider the composed function $g \circ f$. They showed that the conjecture implies a separation of NC^1 from P. In his talk, Or Meir presented a natural step in this direction, which lies between what is known and the original conjecture: he showed that an analogue of the conjecture holds for the composition of a function with a universal relation. The main technical tool is to use information complexity to analyze certain communication problems.

A core theme in circuit complexity is *depth-reduction*: very roughly, these are techniques to reduce the depth of a given circuit without increasing its size too much. The classic work of Valiant, Skyum, Berkowitz and Rackoff shows that any polynomial size arithmetic circuit has an equivalent circuit of polynomial size and $\log^2 n$ depth, where n is the number of input variables. Further impedus was given by Agrawal and Vinay who pushed the depth reduction to constant depth, thereby establishing the *chasm at depth* 4. It states that exponential lower bounds for circuits of depth 4 already give such bounds for general circuits. This was further improved by Koiran and by Tavenas.

Ramprasad Saptharishi gave a slightly different proof of the depth reduction of Tavenas in his talk. Thereby he was able to apply the technique to homogeneous formulas and constant depth formulas.

Chandan Saha presented a very strong result: an exponential lower bound for homogeneous depth-4 circuits that comes close to the chasm-barrier. His techniques also yield exponential lower bounds for certain nonhomogeneous depth-3 circuits. Having the parameters so close to the bounds coming from depth reduction make these results really exciting.

Depth reduction is also an crucial ingredient in Pascal Koirans talk. He presented a new version of the *τ-conjecture* for Newton polygons of bivariate polynomials. The τ-conjecture was originally stated by Shub and Smale:

> *the number of integer roots of a univariate polynomial should be polynomially bounded in the size of the smallest straight-line program computing it.*

Pascal Koiran proposed a new version of the τ-conjecture in his talk:

> *when a bivariate polynomial is expressed as a sum of products of sparse polynomials, the number of edges of its Newton polygon is polynomially bounded in the size of such an expression.*

If this new conjecture is true, then the permanent polynomial cannot be computed by polynomial-size arithmetic circuits.

Spurred by the depth reduction results, we have seen some great work on *Polynomial Identity Testing* (PIT) recently, in particular on depth-3 and depth 4 circuits, and on arithmetic

branching programs. The most ambitious goal here is to come up with a hitting set construction for a specific model. A hitting set is a set of instances such that every non-zero polynomial in the model has a non-root in the set. This solves the PIT problem in the *black box* model.

Rohit Gurjar and Arpita Korwar gave a joint talk on PIT for read-once arithmetic branching programs. They presented a new technique called *basis isolating weight assignment*. These weight assignments yield a hitting set in quasi-polynomial time.

Michael Forbes considered the question whether the hitting set constructions running in quasi-polynomial time can be improved to polynomial time. He showed that in the case of depth-3 powering circuits (sums of powers of linear polynomials) one can obtain a hitting set of size $\mathrm{poly}(s)^{\log\log s}$ for circuits of size s, which is pretty close to resolving the black-box identity testing problem for this class in polynomial time.

Swastik Kopparty showed the computational equivalence of factoring multivariate polynomials and PIT. For both problems we have efficient randomized algorithms. The question whether these algorithms can be derandomized are central in arithmetic complexity. Swastik established that they are equivalent.

Valiant introduced the arithmetic analogue of classes P and NP. Very roughly, the class VP contains all multivariate polynomials that can be computed (non-uniformly) by polynomial-size arithmetic circuits, and the class VNP contains all multivariate polynomials that have coefficients computable by VP-circuits. The question whether VP is different from VNP plays the role of the P-NP question in algebraic complexity theory. Valiant showed that the permanent is complete for VNP. But for VP, only artificially constructed functions were known to be complete. In her talk, Meena Mahajan described several natural complete polynomials for VP, based on the notion of graph homomorphism polynomials.

Eric Allender defined a class called ΛP which is in some sense dual to VP. Over finite fields, VP can be characterized by SAC^1, the class of logarithmic depth, polynomial-size semi-unbounded fan-in circuits (with bounded fan-in multiplication gates and unbounded fan-in addition gates). Eric defined the dual class ΛP in the same way, but with unbounded fan-in multiplication gates and bounded fan-in addition gates. He showed new characterizations of the complexity classes ACC^1 and TC^1 based on ΛP.

Klaus-Joern Lange defined a completeness notion on families of languages, called *densely complete*. He showed that the context-free languages are densely complete in SAC^1 via many-one AC^0-reductions.

Complexity

Ryan Williams once again demonstrated a fruitful interplay between algorithms and complexity. In his famous ACC-paper, he showed how to use fast algorithms for circuit satisfiability to prove lower bounds with respect to the class ACC. In his present talk, Ryan reversed the direction and showed how to exploit techniques from complexity to obtain faster algorithms for the all-pairs shortest paths problem (APSP). He improved the running time from $n^3/\log^2 n$ previously to $n^3/2^{\Omega(\sqrt{\log n})}$. The big question here is whether one can improve the running time to $n^{3-\epsilon}$ for some $\epsilon > 0$. A crucial role in the new algorithm plays the *polynomial method* of Razborov and Smolensky, originally conceived for proving low-depth circuit lower bounds.

Michal Koucký talked on a model of computation he calls *catalytic computation*. In this model, a machine has only limited memory available, but has additionally access to almost unlimited amount of disk space, the *catalytic memory*. This disk is however already full of data. The machine has read-write access to the disk so that it can modify the content of the disk. However, at the end of a computation, the content of the catalytic memory has to be in its original state. The question now is whether the catalytic memory is of any use. Michal showed that a logspace bounded machine with a catalytic memory can do all of nondeterministic logspace. Hence, surprisingly, the catalytic memory really helps, unless L = NL.

Amnon Ta-Shma talked on the problem of *approximating* the eigenvalues of stochastic Hermitian matrices. In an earlier paper he had shown that this is possible in probabilistic logspace in the quantum model of computation, i.e. in BQL. In this talk, Amnon was asking whether this is also possible in probabilistic logspace in the classic world, i.e. in BPL. He showed that how to achieve approximations with *constant* accuracy. To bring the problem into BPL, one would have to approximate the eigenvalues with polynomially small accuracy. This remains open for now.

Venkatesan Guruswami condidered the following promise version of the satisfiability problem: Given a k-SAT instance with the promise that there is an assignment satisfying at least t out of k literals in each clause, can one efficiently find a satisfying assignment? Because 3-SAT is NP-hard, the promise problem is NP-hard for $t \leq k/3$. On the other hand, 2-SAT is efficiently solvable. Extensions of the 2-SAT algorithm show that the promis problem is efficiently solvable for $t \geq k/2$. Venkatesan showed a sharp borderline for the promise problem: it is NP-hard for $t < k/2$. The proof uses part of the PCP-machinery.

Communication Complexity

Amir Yehudayoff talked on communication complexity in the number on the forehead model. He considered the disjointness problem: there are k players, each having a set of numbers from $[n]$. A player can see the numbers of all the other players, but not his own numbers. The task of the payers is to determine, whether there is a number common to all sets. Amir showed a lower bound for the deterministic communication complexity of order $n/4^k$. This is quite amazing since it nearly matches the known upper bound, which is of order $k^2 n/2^k$.

Arkadev Chattopadhyay talked on a communication model, where the inputs are distributed among the vertices of an undirected graph. The vertices coorespond to processors, each processor can send messages only to its neighbors in the graph. Arkadev showed lower bounds on the communication cost for computing certain functions in this model.

Rahul Santhanam considered a communication model called *compression game*. There are two players, Alice and Bob. Alice receives the whole input x and is computationally bounded, by $\mathrm{AC}^0[p]$ in this case, for some prime p. Bob has no information about x and is computationally unbounded. The communication cost of some function f is the number of bits Alice sent to Bob until they agree on the value $f(x)$. Rahul showed a lower bound on the communication complexity of the Mod_q-function, for any prime $q \neq p$.

Coding Theory

Error-correcting codes, particularly those constructed from polynomials, lie at the heart of many significant results in Computational Complexity. Usually, error correcting codes are studied with respect to the Hamming distance. Another model is that of random errors. Amir Shpilka in his talk considered the behaviour of Reed-Muller codes in the Shannon model of random errors. He showed that the rate for Reed-Muller codes with either

low- or high-degree achieves (with high probability) the capacity for the Binary-Erasure-Channel

David Zuckerman talked on the relatively new concept of *non-malleable codes* which was introduced by Dziembowski, Pietrzak, and Wichs in 2010. Informally, a code is non-malleable if the message contained in a modified codeword is either the original message, or a completely unrelated value. Non-malleable codes provide an elegant algorithmic solution to the task of protecting hardware functionalities against "tampering attacks". David showed how to construct efficient non-malleable codes in the so-called C-split-state model that achieve constant rate and exponentially small error.

Game Theory

Steve Fenner considered the following two-player game on a finite partially odered set (poset) S: each player takes turns picking an element x of S and removes all $y > x$ from S. The first one to empty the poset wins. Daniel Grier showed that determining the winner of a poset game is PSPACE-complete. Steve considered the *black-white version* of the game, where each player and each element of S is assigned a color, black or white. Each player is only allowed to remove elements of their own color. He showed that also this black-white version of the poset game is PSPACE-complete. This is the first PSPACE-hardness result known for a purely numerical game. Another interesting result was that the game NimG, a generalization of both Nim and Geography, is polynomial-time solvable when restricted to undirected, bipartite graphs, whereas NimG is known to be PSPACE-complete for general graphs, both directed and undirected.

Bill Gasarch talked on a variant of classical NIM, where there is only one pile of stones and and a given set $\{a_1, a_2, \ldots, a_k\}$ of numbers. A move consists of choosing a number a_i from the set and then removing a_i stones from the pile. The first player who cannot move loses the game. This game has already been well studied. Bill considered an extension of the game where each player starts out with a number of dollars. Now each player has to spend a dollars to remove a stones. He presented some surprising results on the winning conditions for the extended game.

Cryptography

Farid Ablayev generalized classical universal hashing to the quantum setting. He defined the concept of a quantum hash generator and offer a design, which allows one to build a large number of different quantum hash functions. One of the important points here is to use unly few quantum bits. Farid proved that his construction is optimal with respect to the number of qubits needed.

Matthias Krause talked on approaches for designing authentication protocols for ultra-light weight devices as for example RFID chips. He proposed a new approach based on key stream generators as the main building block.

Conclusion

As is evident from the list above, the talks ranged over a broad assortment of subjects with the underlying theme of using algebraic and combinatorial techniques. It was a very fruitful meeting and has hopefully initiated new directions in research. Several participants specifically mentioned that they appreciated the particular focus on a common class of *techniques* (rather than end results) as a unifying theme of the workshop. We look forward to our next meeting!

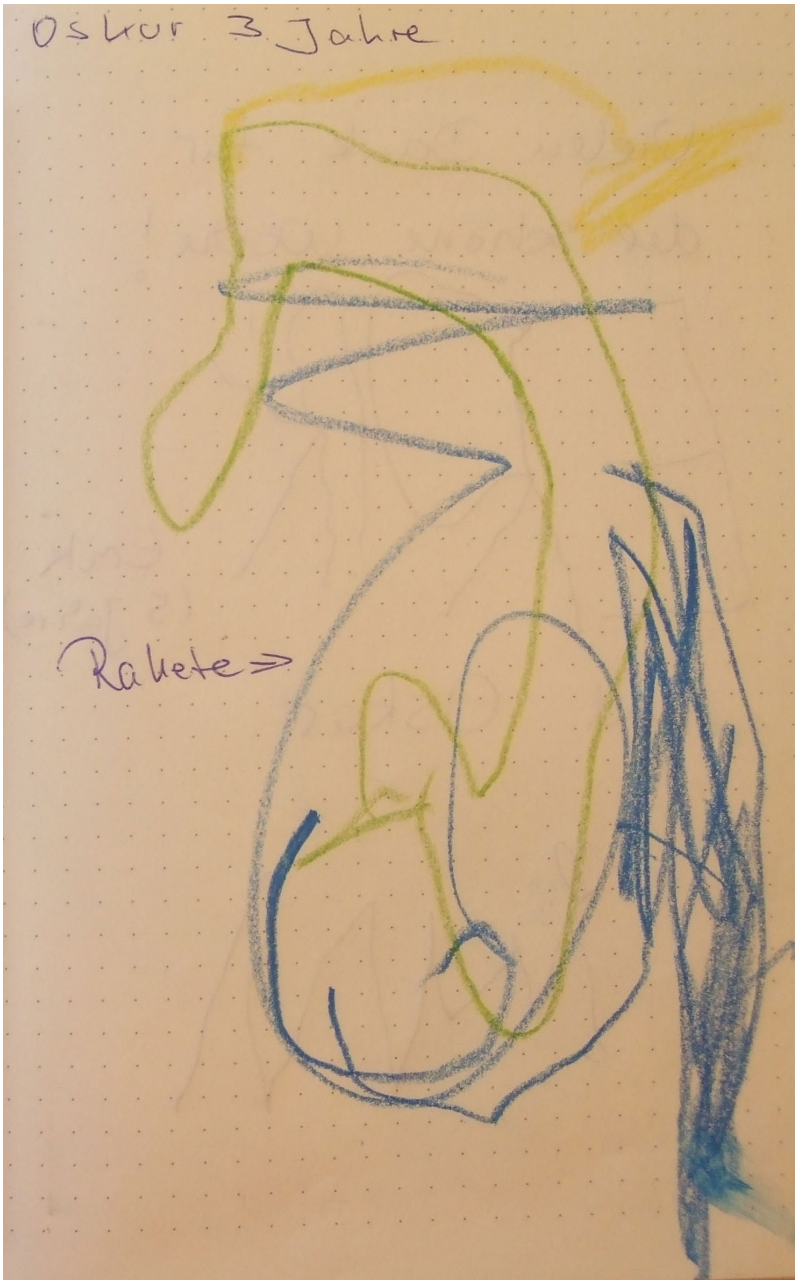

Fig. 4.14
Drawing for the Dagstuhl children's guest book by Oscar (3), son of Dagstuhl Seminar 14132 participants Maike and Kevin Buchin.

4.58 Privacy and Security in an Age of Surveillance

Organizers: Bart Preneel, Phillip Rogaway, Mark D. Ryan, and Peter Y. A. Ryan

Seminar No. 14401

Date: September 28 to October 2, 2014 | Dagstuhl Perspectives Workshop
Full report – DOI: 10.4230/DagRep.4.9.106

Participants: Jacob Appelbaum, Daniel J. Bernstein, Caspar Bowden, Jon Callas, Joseph Cannataci, George Danezis, Pooya Farshim, Joan Feigenbaum, Ian Goldberg, Christian Grothoff, Marit Hansen, Amir Herzberg, Eleni Kosta, Hugo Krawczyk, Susan Landau, Tanja Lange, Kevin S. McCurley, David Naccache, Kenneth G. Paterson, Bart Preneel, Charles D. Raab, Phillip Rogaway, Mark D. Ryan, Peter Y. A. Ryan, Haya Shulman, Vanessa Teague, Vincent Toubiana, Michael Waidner, Dan Wallach

Revelations over the last few years have made clear that the world's intelligence agencies surveil essentially everyone, recording and analyzing who you call, what you do on the web, what you store in the cloud, where you travel, and more. Furthermore, we have learnt that intelligence agencies intentionally subvert security protocols. They tap undersea cables. They install malware on an enormous number targets worldwide. They use active attacks to undermine our network infrastructure. And they use sophisticated analysis tools to profile individuals and groups.

While we still understand relatively little about who is doing what, the documents leaked by Snowden have led to the conclusion that the Five Eyes[19] organizations are going far beyond anything necessary or proportionate for carrying legitimate intelligence activities. ot an equivalent access to documents Governmental assurances of oversight have come to ring hollow, as any oversight to date seems to have been ineffectual, and is perhaps a complete sham.

Can democracy or nonconformity survive if the words and deeds of citizens are to be obsessively observed by governments and their machines? The rise of electronic surveillance thus raises questions of immense significance to modern society. There is an inherent tension. Machine-monitored surveillance of essentially everything people do is now possible. And there are potential economic, political, and safety benefits that power may reap if it can implement effective population-wide surveillance. But there is also a human, social, economic, and political harm that can spring from the very same activity.

The goal of our workshop was to gather together a mix of people with knowledge and expertise in both the legal and technological aspects of privacy and surveillance, to try to understand the landscape that we now live in, and to debate approaches to moving forward. We invited people from a wide range of domains, including members of the intelligence community. All invitees in the intelligence community declined the invitations – in most cases choosing not even to reply. Also, we found that we had more success in getting positive replies from members of the technical community than members of the legal or regulatory communities. Consequently, the makeup of the workshop was not as diverse and balanced as we had hoped. Nonetheless, we felt that we achieved a healthy mix, and there was plenty of lively debate. The issues addressed by this workshop were unusually contentious, and discussions at times were highly animated, even heated.

It is often argued that privacy is not an absolute right. This is true, but this is also true of other rights. The right to freedom must be tempered by the fact that people who are convicted of crimes may forfeit this right for a period. Equally, someone for whom there are sound grounds for suspicion might forfeit some privacy rights. But in any event, any such breaches must be targeted and proportionate and justified by well-founded grounds for suspicion.

An important observation that came up repeatedly in discussions is that privacy is not just an individual right but essential to the health of a democratic society as a whole.

How can society as whole be provided strong assurance that intelligence services are "playing by the rules" while at the same time allowing them sufficient secrecy to fulfill their role? It seems feasible that technical mechanisms can contribute to solving this problem, and indeed a number of presentations addressed aspects of it. One might imagine that something analogous to the notion of zero-knowledge proofs might help demonstrate that intelligence agencies are following appropriate rules while not revealing details of those activities. Another possibility that was proposed is to make the amount of surveillance public in a verifiable fashion but without revealing the targets. Thus one

[19] This term is used to indicate Australia, Canada, UK, USA, and New Zealand.

might imagine that a specified limit be placed on the proportion of traffic available to intelligence services. The effect would be to force the agencies to be correspondingly selective in their choice of targets.

The crypto and security community should invest a substantial effort to make all layers of the internet and our devices more secure and to strengthen the level of privacy offered. This may create a natural barrier to mass surveillance and will also bring a more robust network infrastructure to a society that is increasingly reliant on it for critical services. Such a development may eventually increase the cost for targeted surveillance, but there is no indication that this would become prohibitive.

As is traditional for Dagstuhl, we started with a round table of quick introductions from the participants, including brief statements of what they hoped to get out of the workshop. We then had an open discussion on the goals of the workshop and of how best to organise the workshop to achieve these goals. It was decided to structure discussions into three strands:

- Principles
- Research directions
- Strategy

The outcomes of these discussions are detailed in a separate "Manifesto" document. The workshop was then structured into a number of plenary sessions alternating with breakouts into the three strands. The plenary sessions were made up of presentations from participants and feedback from the breakouts followed by discussion.

The problems addressed in this workshop are immensely challenging, and carry vast implications for society as a whole. It would not be reasonable to expect a small group of people – and a group not particularly representative of society as a whole – to produce solutions in the course of four days. Our goal was to gain some understanding of guiding principles and ways forward.

4.59 Resilience in Exascale Computing

Organizers: Hermann Härtig, Satoshi Matsuoka, Frank Mueller, Alexander Reinefeld

Seminar No. 14402

Date: September 28 to October 1, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.9.124

Participants: Dorian C. Arnold, Amnon Barak, Leonardo Bautista-Gomez, George Bosilca, Zizhong Chen, Andrew A. Chien, Nathan DeBardeleben, Björn Döbel, James J. Elliott, Christian Engelmann, Mattan Erez, Hermann Härtig, Torsten Hoefler, Larry Kaplan, Dieter Kranzlmüller, Matthias Lieber, Naoya Maruyama, Satoshi Matsuoka, Frank Mueller, Alexander Reinefeld, Yves Robert, Robert B. Ross, Kento Sato, Thorsten Schütt, Martin Schulz, Vilas Sridharan, Thomas Steinke, Jeffrey Vetter, Abhinav Vishnu, Gerhard Wellein, Felix Wolf

Motivation

The upcoming transition[20] from petascale to exascale computers requires the development of radically new methods of computing. Massive parallelism, delivered by manycore processors and their assembly to systems beyond 10^7 processing units will open the way to extreme computing with more than 10^{18} floating point operations per second. The large number of functional components (computing cores, memory chips, network interfaces) will greatly increase the probability of partial failures. Already today, each of the four fastest supercomputers in the TOP500 list[21] comprises more than half a million CPU cores, and this tendency towards massive parallelism is expected to accelerate in the future. In such large and complex systems, component failures are the norm rather than an exception. Applications must be able to handle dynamic reconfigurations during runtime and system software is needed to provide fault tolerance (FT) at a system level. For example, Jaguar reportedly experience 20 faults per hour in production mode[22], some of which could be mitigated while others could not.

To prevent valuable computation to be lost due to failures, checkpoint/restart (C/R) has become a requirement for long running jobs. However, current C/R mechanisms are insufficient, because the communication channels between the main memory and the parallel file system are far too slow to allow to save (and restore) a complete memory dump to disk farms. As an alternative, the memory of neighboring compute nodes may be used to keep partial checkpoints, but then erasure coding must be used to prevent against the loss of data in case of single node failures. To make things worse, precious communication bandwidth is needed for writing/reading checkpoints, which slows down the application. Techniques for data compression or application-specific checkpointing (with a reduced memory footprint) were proposed as a solution, but they only alleviate the problem by a certain extent.

We assume exascale hardware architectures to consist of a heterogeneous set of computational units (ranging from general-purpose CPUs to specialized units such as today's GPUs), memory chips (RAM, flash, phase-change memory), and various kinds of interconnects. The operating system and its load balancing mechanisms need to adapt to the hardware's properties as well as to workload characteristics. With the co-existence of legacy applications and new applications, it can be assumed that exascale systems must be capable of executing a broad range of parallel programming paradigms like MPI, OpenMP, PGAS, or MapReduce. These will not always and in every case require the functionality of a fully fledged operating system. We furthermore expect applications to become more complex and dynamic. Hence, developers cannot be expected to continuously handle load balancing and reliability. It is the operating system's task to find a sweet spot that on the one hand provides generic means for load management and checkpointing, while on the other hand allowing application developers full control over the performance-relevant functionality if required.

[20] IDC top ten market prediction no. 2: *The Global Petascale/Exascale Race Will Keep Shifting the Market Toward Larger Systems*, IDC, March 2013.
[21] http://www.top500.org
[22] A. Geist, "What is the Monster in the Closet?", August 2011, Invited Talk at Workshop on Architectures I: Exascale and Beyond: Gaps in Research, Gaps in our Thinking

Objectives and Expected Results

The objective of this seminar is to bring together researchers and developers with a background on HPC system software (OS, network, storage, management tools) to discuss medium to long-term approaches towards resilience in exascale computers. Two concrete outcomes are (a) outlines for alternatives for resilience at extreme scale with trade-offs and dependencies on hardware/technology advances and (b) initiation of a standardization process for a resilience API. The latter is driven by current trends of resilience libraries to let users specify important data regions required for tolerating faults and for potential recovery. Berkeley Lab's BLCR, Livermore's SCR and Capello's FTI feature such region specification in their APIs, and so do may in-house application-specific solutions. A standardized resilience API would allow application programmers to be agnostic of future underlying resilience mechanisms and policies so that resilience libraries can be exchanged at will (and might even become inter-operable). The focus of solutions is on the practical system side and should reach beyond currently established solutions. Examples of areas of interest are:

- What is the "smallest denominator" that defines a resilience API? How can the standardization of a resilience API be realized?
- How can reactive FT schemes that respond to failures be enhanced to reduce system overhead, ensure progress in computation and sustain ever shorter MTBFs?
- How should low-energy and/or persistent memory be included on nodes for checkpointing (for example PCM) and used by applications and the OS?
- Can a significant number of faults be predicted with exact locations ahead of time so that proactive FT may provide complementary capabilities to move computation away from nodes that are about to fail?
- Can message logging, incremental checkpointing and similar techniques contribute to lower checkpointing overhead?
- Is redundant execution a viable alternative at exascale? How can partly redundant execution contribute to increased resilience in exascale algorithms?
- Can algorithm-based fault tolerance be generalized to entire classes of algorithms? Can results continuously be checked?
- What is the impact of silent data corruption (SDC) on HPC computing today? Which solvers can tolerate SDCs, which ones need to be enhanced (and how)?
- How do current/novel network architectures interact with the OS (e. g., how does migration interact with RDMA)?
- How can execution jitter be reduced or tolerated on exascale systems, particularly in the presence of failures?
- Can an interface be designed that allows the application to give "hints" to the OS in terms of execution steering for resilience handling? How does this approach interact with scalability mechanisms and policies, e. g., load balancing, and with programming models, e. g., to define fault handlers?
- Do distributed communication protocols offer better resilience? How do they support coordination between node-local and inter-node scheduling?
- Does "dark silicon" offer new opportunities for resilience?
- How can I/O on exascale be efficient and resilient (e. g., in situ analysis of simulation results)?

As a result of the seminar, we expect that this list of objectives will be refined, extended, and approaches to address each of these problems will be formulated, We anticipate that participants engage in increased coordination and collaboration within the currently (mostly) separate communities of HPC system software and application development.

Furthermore, the standardization process will be kicked off. One challenge is to find the most promising context for standardization. Current HPC-related standards (MPI, OpenMP, OpenACC) do not seem suitable since resilience cuts across concrete runtime environments and may also extend beyond HPC to Clouds and data centers involving industry participants from these area (in future standardization meetings beyond the scope of this meeting).

Overall, the objective of the workshop is to spark research and standardization activities in a coordinated manner that can pave the way for tomorrow's exascale computers to the benefit of the application developers. Thus we expect not only HPC system developers to benefit by the seminar but also the community of scientific computing at large, well beyond computer science. Due to the wide range of participants (researchers and industry practitioners from the U.S., Europe, and Asia), forthcoming research work may significantly help enhance FT properties of exascale systems, and technology transfer is likely to also reach general-purpose computing with many-core parallelism and server-style computing. Specifically, the work should set the seeds for increased collaborations between institutes in Europe and the U.S./Asia.

Relation to Previous Dagstuhl Seminars

Two of the organizers, Frank Mueller and Alexander Reinefeld, previously co-organized a Dagstuhl Seminar on *Fault Tolerance in High-Performance Computing and Grids* in 2009. It provided a forum for exchanging research ideas on FT in high-performance computing and grid computing community. In the meantime, the state-of-the-art greatly advanced and it became clear, that exascale computing will not be possible without adequate means for resilience. Hence, the seminar was more concrete in that the pressing problems of FT for exascale computing and standardization must be tackled and solved with the joint forces of system researchers and developers.

The seminar also builds on the Dagstuhl Perspective Workshop 12212 *Co-Design of Systems and Applications for Exascale*, which also relates to the DFG-funded project FFMK (http://ffmk.tudos.org/, "A Fast and Fault-tolerant Microkernel-based System for Exascale Computing", DFG priority program 1648). Compared to the perspective workshop, our seminar was much more focused on a single, pressing topic of exascale computing, namely resilience.

4.60 Constraints, Optimization and Data

Organizers: Luc De Raedt, Siegfried Nijssen, Barry O'Sullivan, and Michele Sebag

Seminar No. 14411

Date: October 6–10, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.10.1

Participants: Hendrik Blockeel, Jean-François Boulicaut, Ken Brown, Bruno Crémilleux, James Cussens, Krzysztof Czarnecki, Thi-Bich-Hanh Dao, Ian Davidson, Luc De Raedt, Marc Denecker, Yves Deville, Alan Frisch, Randy Goebel, Valerio Grossi, Tias Guns, Holger H. Hoos, Frank Hutter, Kristian Kersting, Lars Kotthoff, Pauli Miettinen, Mirco Nanni, Benjamin Negrevergne, Siegfried Nijssen, Barry O'Sullivan, Andrea Passerini, Francesca Rossi, Lakhdar Sais, Vijay A. Saraswat, Michele Sebag, Arno Siebes, Guido Tack, Yuzuru Tanaka, Joaquin Vanschoren, Christel Vrain, Toby Walsh

Constraint programming and optimization (CPO) have recently received considerable attention from the fields of machine learning and data mining (MLDM). On the one hand, the hypotheses and patterns that one seeks to discover in MLDM can be specified in terms of constraints (e. g. labels in the case of supervised learning, preferences in the case of learning to rank, must-link and cannot-link in the case of unsupervised learning, coverage and lift in the case of data mining). On the other hand, powerful constraint programming solvers have been developed. If MLDM users express their requirements in terms of constraints they can delegate the MLDM process to such highly efficient solvers.

Conversely, CPO can benefit from integrating learning and mining functionalities in a number of ways. For example, formulating a real-world problem in terms of constraints requires significant expertise in the problem domain. Also, selecting the most appropriate constraints, in terms of constraint solving efficiency, requires considerable expertise in the CPO domain. In other words, experience plays a major role in successfully applying CPO technology.

In addition, both CPO and MLDM share a common challenge associated with tuning their respective methods, specifically determining the best parameters to chose for an algorithm depending on the task at hand. A typical performance metric in machine learning is the predictive accuracy of a hypotheses, while in CPO it might be search cost or solution quality.

This seminar built upon the 2011 *Constraint Programming meets Machine Learning and Data Mining*[23] and the 2014 *Preference learning*[24] seminars. Its goal was to identify the key challenges and opportunities at the crossroads of CPO and MLDM. The interests of the participants included the following:

- Problem formulation and modelling: constraint-based modelling; preference formalisms; loss functions in ML; modelling and exploiting background knowledge; structured properties (e. g. preserving spatio-temporal structures).
- Improvement of algorithms / platforms in the areas of algorithm selection, algorithm configuration, and/or algorithm scheduling, particularly with respect to parallel execution.
- Specification and reasoning about goals and optimization criteria: modelling preferences and integrating with human expertise (exploiting the "human in the loop") to converge on high quality outcomes.
- Additional functionalities such as the use of visualization and explanation.
- Algorithmic scalability.
- Approximate reasoning, reasoning under uncertainty, and incorporating probability.

The seminar was organized into seven sessions: frameworks and languages; algorithm configuration; constraints in pattern mining; learning constraints; machine learning with constraints; applications; and demonstrations. The demonstrations presented at the seminar were by:

- Guido Tack – MiniZinc (see http://mininzinc.org);
- Joaquin Vanschoren – OpenML (see http://openml.org);
- Tias Guns – MiningZinc (see http://dtai.cs.kuleuven.be/CP4IM/miningzinc);
- Bruno Crémilleux – software for the calculation of Sky Pattern Cubes;
- Marc Denecker – IDP (see http://dtai.cs.kuleuven.be/krr/software/idp);
- Holger Hoos – algorithm selection and portfolio software;

[23] http://www.dagstuhl.de/11201
[24] http://www.dagstuhl.de/14101

- Luc De Raedt – ProbLog (see http://dtai.cs.kuleuven.be/problog/).

The seminar also had five working groups on:

- Declarative Languages for Machine Learning and Data Mining;
- Learning and Optimization with the Human in the Loop;
- Meta-Algorithmic Techniques;
- Big Data;
- Towards Killer Applications.

4.61 Globalizing Domain-Specific Languages

Organizers: Betty H. C. Cheng, Benoit Combemale, Robert B. France, Jean-Marc Jezequel, and Bernhard Rumpe

Seminar No. 14412

Date: October 5–10, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.10.32

Participants: Colin Atkinson, Cedric Brun, Barrett Bryant, Benoit Caillaud, Betty H. C. Cheng, Tony Clark, Siobhán Clarke, Benoit Combemale, Julien Deantoni, Thomas Degueule, Robert B. France, Ulrich Frank, Jean-Marc Jezequel, Gabor Karsai, Ralf Lämmel, Marjan Mernik, Pieter J. Mosterman, Oscar M. Nierstrasz, Bernhard Rumpe, Martin Schindler, Friedrich Steimann, Eugene Syriani, Janos Sztipanovits, Juha-Pekka Tolvanen, Antonio Vallecillo, Mark van den Brand, Markus Völter

Model Driven Engineering (MDE) aims to reduce the accidental complexity associated with developing complex software-intensive systems, through the development of technologies that enable developers to systematically create, evolve, analyse, and transform various forms of abstract system models.

Current MDE language workbenches, in both academia and industry, support the development of Domain-Specific Modeling Languages (DSMLs) that can be used to create models that play pivotal roles in different development phases. Language workbenches such as EMF, Metaedit+ or MPS support the specification of the abstract syntax, concrete syntax and the static and dynamic semantics of a DSML. These workbenches aim to address the needs of DSML developers in a variety of application domains.

The development of modern complex software-intensive systems often involves the use of multiple DSMLs that capture different system aspects. In addition, models of the system aspects are seldom manipulated independently of each other. System engineers are thus faced with the difficult task of relating information presented in different models. Current DSML development workbenches provide good support for developing independent DSMLs, but provide little or no support for integrated use of multiple DSMLs. The lack of support for explicitly relating concepts expressed in different DSMLs (incl., syntax and semantics) makes it very difficult for developers to reason about information spread across models describing different system aspects.

Supporting coordinated use of DSMLs leads to what we call the globalization of modeling languages, that is, the use of multiple modeling languages to support coordinated development of diverse aspects of a system.

Discussions during the seminar will focus on how multiple heterogeneous modeling languages (or DSMLs) will need to be related to determine how different aspects of a system influence each other. We have identified three forms of relationships among DSMLs that can be used as a starting point for discussions: interoperability, collaboration, and composition. These forms of language integration will need to address challenging issues that arise from the heterogeneity of modeling languages. Relationships among the languages will need to be explicitly defined in a form that corresponding tools can use to realize the desired interactions. Requirements for tool manipulation is thus another topic that will be discussed in the seminar.

The goal of the seminar was to develop a research program that broadens the current DSML research focus beyond the development of independent DSMLs to one that provides support for globalized DSMLs. In the globalized DSMLs vision, integrated DSMLs provide the means for teams working on systems that span many specialized domains and concerns to determine how their work on a particular aspect influences work on other aspects.

Working Groups

In the seminar we started the following four working groups which are producing results during the workshop and compiling them into a state-of-the-art report afterwards:

Group 1a Motivating Use Cases for the Globalization of DSLs Definition of the main scenarios motivating the globalization of DSLs

Group 1b Conceptual Model of the Globalization of Domain-Specific Languages Definition of the common vocabulary and foundations of the globalization of DSLs

Group 2 Globalized Domain Specific Language Engineering Challenges of the globalization of DSLs from the language designer point of view

Group 3 Domain Globalization: Using Languages to Support Technical and Social Coordination Challenges of the globalization of DSLs from the language user point of view

Fig. 4.15
Maria Satter – Untitled. Part of the Dagstuhl art collection and donated by the participants in Dagstuhl Seminar 03111.

4.62 Optimal Algorithms and Proofs

Organizers: Olaf Beyersdorff, Edward A. Hirsch, Jan Krajíček, and Rahul Santhanam

Seminar No. 14421

Date: October 12–17, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.10.51

Participants: Per Austrin, Olaf Beyersdorff, Ilario Bonacina, Samuel R. Buss, Igor Carboni Oliveira, Ruiwen Chen, Yijia Chen, Leroy Chew, Andrew Drucker, Susanna Figueiredo de Rezende, Jörg Flum, Nicola Galesi, Michal Garlik, Christian Glasser, Andreas Goerdt, Johan Hastad, Edward A. Hirsch, Dmitry Itsykson, Mikoláš Janota, Emil Jerabek, Alexander Knop, Johannes Köbler, Jan Krajíček, Alexander S. Kulikov, Massimo Lauria, Barnaby Martin, Jochen Messner, Hunter Monroe, Moritz Müller, Sebastian Mueller, Jakob Nordström, Jan Pich, Pavel Pudlák, Benjamin Rossman, Zenon Sadowski, Rahul Santhanam, Alan Selman, Alexander Smal, Dmitry Sokolov, David Steurer, Jacobo Torán

General Introduction to the Topic

The notion of optimality plays a major role in theoretical computer science. Given a computational problem, does there exist a "fastest" algorithm for it? Which proof system yields the shortest proofs of propositional tautologies? Is there a single distribution which can be used to inductively infer any computable sequence? Given a class of optimization problems, is there a single algorithm which always gives the best efficient approximation to the solution? Each of these questions is a foundational one in its area – the first in computational complexity, the second in proof complexity, the third in computational learning theory, and the last in the theory of approximation.

Consider, as an example, the Boolean Satisfiability (SAT) search problem, which asks, given a Boolean formula, for a satisfying assignment to the formula. Since SAT is NP-complete, being able to tell whether the fastest algorithm for SAT runs in polynomial time would imply a solution to the notoriously hard NP vs P problem, which is far beyond the state of our current knowledge. However, the possibility remains that we can define an *optimal* algorithm which we can guarantee to be essentially the fastest on every instance, even if we cannot rigorously analyze the algorithm. In a seminal paper, Leonid Levin (1973) proved that every NP search problem, and in particular SAT, has an optimal algorithm. It is still unknown whether every decision problem in NP has an optimal algorithm.

In general, given a class of computational artefacts (algorithms/proof systems/distributions) and performance measures for each artefact in the class, we say that an artefact is optimal if it matches the performance of every other artefact in every case. The main questions about optimality is: for which classes of artefacts and under which assumptions do they exist? In case they do exist, how well do they match the performance of other artefacts in the class? How is the existence of optimal artefacts related to other fundamental theoretical questions, such as complexity lower bounds, efficient learnability or approximability?

There have been a number of important recent results about optimality in various computational settings. Prime examples include optimal proof systems and acceptors under advice or in heuristic settings, surprising relations of optimal proof systems to descriptive complexity and parameterized complexity, hierarchy results in various computational settings, and optimal approximation algorithms for constraint satisfaction problems.

Organisation of the Seminar and Activities

The seminar brought together 41 researchers from different areas of computer science and mathematics such as computational complexity, proof complexity, logic, and approximations with complementary expertise, but common interest in different notions of optimality. The participants consisted of both senior and junior researchers, including a number of postdocs and a few advanced graduate students.

Participants were invited to present their work and to communicate state-of-the-art advances. Twenty-two talks of various lengths were given over the five-day workshop. Survey talks of 60 minutes were scheduled prior to workshop, covering the three main areas of computational complexity, proof complexity, and approximations. Most of the remaining slots were filled as the workshop commenced. In addition, during two spontaneously organised open problem sessions – one at the very start and the second, longer one near the end of the workshop – the participants posed a number of open problems across the different disciplines covered by the seminar. The organisers considered it important to leave ample free time for discussion.

Three tutorial talks were scheduled during the first two days in order to establish a common background for the different

communities from computational complexity, proof complexity, logic, and approximation that came together for the workshop. The presenters and topics were:

- David Steurer: Survey on Approximations and Optimality
- Olaf Beyersdorff: Optimal Proof Systems – a Survey
- Rahul Santhanam: Hierarchies and Lower Bounds via Optimality – a Survey

The other 19 talks covered a broad range of topics from logic, computational complexity and proof complexity.

The organisers think that the seminar fulfilled their original high goals: most talks were a great success and many participants reported about the inspiring seminar atmosphere, fruitful interactions, and a generally positive experience. The organisers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and excellent support in the arrangement of a very successful and productive event.

4.63 Modeling, Verification, and Control of Complex Systems for Energy Networks

Organizers: Alessandro Abate, Martin Fränzle, Ian Hiskens, and Martin Střelec

Seminar No. 14441

Date: October 26–31, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.10.69

Participants: Alessandro Abate, Erika Ábrahám, Christel Baier, Calin A. Belta, Marc Bouissou, Ana Busic, Claudio De Persis, Florian Dörfler, Martin Fränzle, Sicun Gao, Boudewijn Haverkort, Holger Hermanns, Ian Hiskens, Ondřej Holub, Gabriela Hug-Glanzmann, George B. Huitema, Anak Agung Julius, Maryam Kamgarpour, Stamatis Karnouskos, Joost-Pieter Katoen, David Kleinhans, Jan Krčál, Rom Langerak, Mats Larsson, Sebastian Lehnhoff, Bernard Lesieutre, Roland P. Malhame, Kostas Margellos, Johanna Mathieu, Patrick Panciatici, Maria Prandini, Stefan Ratschan, Anne Remke, Sandip Roy, Jeremy Sproston, Martin Střelec, Ufuk Topcu, Arjan van der Schaft, Lijun Zhang

The seminar has been focused on a number of selected topics from energy networks, with an emphasis on power systems that have great societal and economical relevance and impact. These represent systems of considerable engineering interest, since:

- they can be large-scale and can involve numbers of various devices interconnected in a complex manner.
- they are heterogeneous, that is they can be naturally modelled through a combination of continuous dynamical elements (to capture the evolution of quantities such as voltages, frequencies and generation output) and discrete dynamical components (to capture changes in the network topology, controller logic, state of breakers, isolation devices, transformer taps, etc.).
- they involve substantial stochastic components. Sources of uncertainty traditionally considered in power networks include hardware faults and unforeseen events, as well as stochasticity arising from continuous processes, particularly power demand. Furthermore, the increasing availability of renewable energy sources (e. g. photovoltaic panels, wind turbines, etc.) implies that uncertainty (for example, uncertainty in weather forecasts or cloud cover) also enters at the power supply side.
- some variables are only partially observable due to absence of real-time sensing circuitry in large parts of the existing power distribution network.

Reasonable and accurate analysis of future power networks needs models that seamlessly integrate behavioural patterns like complex interaction of continuous electrical phenomena (e. g. power flows) related to connected devices, discrete events caused by switching behaviour in circuitry, commitment of supplies and loads or by decisions of market participants, and the inherently stochastic behaviour of volatile supplies, demands and market prices.

In summary, the aim of the seminar has been to survey existing and explore novel formal frameworks for modelling, analysis and control of complex, large scale systems, with emphasis on applications in power networks. The seminar has hosted researchers and practitioners working on energy network application domains, in order to import related techniques for the study of energy grids in general, their analysis and energy management, which consists in control, coordination and dispatch of multiple generation, consumption and storage devices connected to the grid. Interactions among scientists and professionals from the heterogeneous research and application fields focused on power networks has highlighted opportunities for further research concerning expressiveness of models and scalability of the methods, as well as point to related efforts in the power network community.

General comments

The Seminar has run over the last week of October 2014 (27 to 31), has been well attended throughout the week, with about 40 participants. It has featured a fully packed program made up of presentations (at least 30), sustained discussions, and breakout sessions on three different topics. A final discussion session has concluded the proceedings of this event.

While the presence from academia has been preponderant, we have also been happy to see a number of active participants from the industry. The attendants expertise has been quite diverse. Academic participants have come with backgrounds in verification, control, and power systems. Alongside the participated and very open discussions, the seminar has additionally featured a hike and a dinner at a local restaurant.

Program

Talks have been categorised within the following clusters: Theory and Tools from Control; Theory and Tools from Verification; Topics in Power Networks; Smart/Micro Grids and Buildings.

Beyond these clusters, we have tried to diversify the program in order to optimally engage the audience. Discussions have been fostered via an afternoon breakout session, organised on Tuesday, the social activities on Wednesday afternoon, and the final session on Friday in the late morning.

There have been three breakout sessions, focusing respectively on

- modelling issues in energy/power systems;
- simulation issues in energy/power systems;
- demand response: control and verification.

The topics elaborated during the sessions are discussed in the full report, which also reports the notes that have come out of the discussions.

4.64 Optimality and Tight Results in Parameterized Complexity

Organizers: Stefan Kratsch, Daniel Lokshtanov, Dániel Marx and Peter Rossmanith

Seminar No. 14451

Date: November 2–7, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.1

Participants: Jakob Lykke Andersen, Carsten Baldauf, Wolfgang Banzhaf, Nikolaj Bjørner, Sebastian Böcker, David Dill, Peter Dittrich, Andreas Dress, Rolf Fagerberg, Harold Fellermann, Christoph Flamm, Martin Hanczyc, Marc Hellmuth, Wim Hordijk, Takashi Ikegami, Ádám Kun, Barry McMullin, Daniel Merkle, Markus E. Nebel, Sijbren Otto, Adolfo Piperno, D. Eric Smith, Pietro Speroni di Fenizio, Peter F. Stadler, Darko Stefanovic, Ralf Steuer, Klaus-Peter Zauner

While many seemingly hard computational problems can be solved satisfactorily in practice, classical complexity dictates that they are in fact intractable (NP-hard) in general. This is an unsatisfactory situation since one would desire a more productive interplay between more heuristic practical results and theoretically proven theorems.

Parameterized complexity analyzes the complexity in finer detail by considering additional problem parameters beyond the input size and expresses the efficiency of the algorithms in terms of these parameters. In this framework, many NP-hard problems have been shown to be (fixed-parameter) tractable when certain structural parameters of the inputs are bounded. In the past two decades, there has been tremendous progress in understanding which problems are fixed-parameter tractable and which problems are not (under standard complexity assumptions).

In recent years, the field of parameterized complexity seems to have evolved beyond merely classifying problems as fixed-parameter tractable or not. The focus shifted to understanding how close the algorithmic results are to the "best possible" algorithm for the problem. Thanks to significant recent advances on both algorithms (upper bounds) and complexity (lower bounds), we have now a tight understanding of many problems and many algorithmic results can be now proven optimal under reasonable assumptions. Moreover, it turns out that the search for optimality can be formulated with respect to different aspects of parameterized complexity and each such aspect gives a separate challenging but doable research direction. One can consider the optimality of algorithms for parameterized problems (either fixed-parameter tractable or not), the optimality of preprocessing algorithms, and the optimality of algorithms with respect to the generality of the problem being solved. The goal of the seminar was to bring together experts in the area of parameterized complexity and algorithms, highlight these research directions and the relevant recent results, and discuss future research topics.

The scientific program of the seminar consisted of 25 talks. Among these there were five 60 minute tutorials on the core topics of the seminar: Marek Cygan and Michał Pilipczuk ("Exponential Time Hypothesis, Part 1+2") covered the Exponential Time Hypothesis (ETH), focussing on techniques for proving tight runtime lower bounds under ETH. Daniel Lokshtanov ("The Strong Exponential Time Hypothesis") introduced Strong ETH as well as related lower bound techniques, and Virginia Vassilevska Williams ("Implications of SETH for polynomial time problems") gave an overview of tight lower bounds for efficiently solvable problems under Strong ETH. Finally, Dániel Marx ("Every Graph is easy or hard") covered the topic of dichotomy theorems for graph problems. Throughout, the tutorials were well-received both as a means of introduction to the topics but also as a convenient way of catching up on very recent results pertaining to the seminar. Furthermore, with most tutorials being held on Monday and Tuesday morning this set a productive atmosphere for tackling open problems regarding tight parameterized complexity results. A further 60 minute contributed talk by Saket Saurabh discussed the recent breakthrough result of fixed-parameter tractability of Graph Isomorphism with respect to treewidth. 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. An open problem session was organized on Monday evening. Notes on the presented problems can be found in the full report.

4.65 Algorithmic Cheminformatics

Organizers: Wolfgang Banzhaf, Christoph Flamm, Daniel Merkle, and Peter F. Stadler

Seminar No. 14452

Date: November 2–7, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.22

Participants: Jakob Lykke Andersen, Carsten Baldauf, Wolfgang Banzhaf, Nikolaj Bjørner, Sebastian Böcker, David Dill, Peter Dittrich, Andreas Dress, Rolf Fagerberg, Harold Fellermann, Christoph Flamm, Martin Hanczyc, Marc Hellmuth, Wim Hordijk, Takashi Ikegami, Ádám Kun, Barry McMullin, Daniel Merkle, Markus E. Nebel, Sijbren Otto, Adolfo Piperno, D. Eric Smith, Pietro Speroni di Fenizio, Peter F. Stadler, Darko Stefanovic, Ralf Steuer, Klaus-Peter Zauner

Dagstuhl Seminar 14452 "Algorithmic Cheminformatics" was organized to intensify the interactions between chemistry and computer science. While the thriving field of bioinformatics/computational biology is a success story of a lively and extensive inter- and trans-disciplinary collaboration between life sciences and computer science, this is much less so in cheminformatics. After a quick raise of a plethora of computational approaches for chemical problems in the 1960–1970s, the field mainly settled down on machine learning approaches in the late 1990s. Over last two decades, computer science plays a comparably marginal role in chemistry research and education.

This is a puzzling state of affairs as chemistry, and in particular the emerging field of systems chemistry, has to offer a wide range of non-trivial computational problems that are very different from those in the well-established areas of quantum chemistry, molecular dynamics, or physical chemistry, for which physics-style models and numeric mathematics have been established as the methods of choice. In particular, complex chemical networks capable of algorithmic self-assembly under far-from-equilibrium conditions, seem to possess a deep connection to the theory of computation, information recoding and compiler theory. Dagstuhl Seminar 14452 therefore specifically aimed to establish the connection between theoretical computer science, graph theory and related fields of discrete mathematics, and complexity theory on the one hand and chemistry on the other hand. Several key areas where covered by one or more presentation and extensive discussions among the participants. Topic ranged from formalizing chemical transformations, autocatalytic molecular systems, and the design of chemical experiments, via model checking and key graph algorithm, to chemical information technology and models for the origin of life. Dagstuhl Seminar 14452 successfully brought together wet-lab chemists with theoretical computer scientists and researchers with a focus on bioinformatics and initiated an, as we feel, very fruitful frist step towards cross-boundary research.

4.66 High-performance Graph Algorithms and Applications in Computational Science

Organizers: Ulrich Carsten Meyer, Henning Meyerhenke, Ali Pinar, and Ilya Safro

Seminar No. 14461

Date: November 10–14, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.40

Participants: Deepak Ajwani, Elisabetta Bergamini, Rob Bisseling, Erik Boman, Christian Brugger, Aydin Buluc, Ümit V. Çatalyürek, Deepayan Chakrabarti, Tiago de Paula Peixoto, Yann Disser, John Feo, Enver Kayaaslan, Dominique LaSalle, Andrew Lumsdaine, Kamesh Madduri, Aleksander Madry, Fredrik Manne, Ulrich Carsten Meyer, Friedhelm Meyer auf der Heide, Henning Meyerhenke, Benjamin A. Miller, Petra Mutzel, Braxton Osting, Srinivasan Parthasarathy, Francois Pellegrini, Ali Pinar, Alex Pothen, Ilya Safro, Peter Sanders, Christian Schulz, Anand Srivastav, Christian Staudt, Veronika Strnadova, Sivan Toledo, Jesper Larsson Träff, Bora Ucar, Panayot S. Vassilevski, David Veith, Katharina A. Zweig

Many presentations in this Dagstuhl seminar emphasized recent trends regarding typical inputs and their effect on graph algorithm development for practical purposes. One can divide these presentations into four categories: (i) Traditional graph (or matrix) problems in new scenarios, (ii) graph analytics algorithms of various sorts, (iii) parallel computing aspects such as tools, computational models, load balancing, and communication; and finally (iv) emerging high-performance application and hardware trends. The following four paragraphs give a brief overview over the talks presented in each of these categories.

Pothen discussed different matching problems and how the emergence of complex networks have changed various matching algorithms recently. Road networks, in turn, are by no means complex and the traditional Dijkstra algorithm solves queries on continental instances in few seconds. Yet, for more challenging scenarios, for example millions of queries per second on webservers or multiple optimization critieria, more elaborate solutions are necessary, as presented by Sanders. Toledo addressed the importance of communication efficiency on large-scale parallel systems for traditional numerical problems such as LU decomposition. A similar numerical topic was the solution of Laplacian linear systems, for which new combinatorial solvers and related techniques from the theory community were presented and discussed by Madry and by Toledo. Furthermore, Boman and Toledo initiated a tangible plan for a scientific competition on solvers for this class of linear systems.

The analytics algorithms part experienced a number of talks on graph clustering and community detection, which means the identification of natural vertex groups in graphs. Several very fast algorithms and their implementation were discussed and compared. Centralities are used for finding important (but in general unrelated) vertices or edges in a graph. Çatalyürek showed how to exploit parallelism in centrality algorithms to speed them up in different hardware settings, including accelerators. Bergamini, in turn, used approximation to obtain a speedup in dynamic graphs. Many other analytics tasks and algorithms were discussed, including anomaly detection presented by Miller and label inference by Chakrabarti, who both focused on techniques for very large graphs. Graph size was also a motivation for sparsification as discussed by Parthasarathy, either to save space or running time (or both) in later stages of an algorithmic pipeline.

Parallelism was the common theme in the third category. Here we summarize algorithmic techniques such as load balancing by graph partitioning, computational models as well as tools and middleware. Several speakers outlined challenges and/or algorithmic solutions in graph partitioning, in particular for complex networks or massively parallel systems. It became also clear that the development of graph algorithms for massive inputs benefits from suitable computational models. An example is the parallel external memory model for which Meyer as well as Veith showed algorithmic solutions. Another prerequisite for efficient graph algorithms in practice is tool support, including building block standards (proposed by Buluc) and communication middleware (presented by Lumsdaine). The pros and cons of different tools were discussed in an animated manner with the co-located Dagstuhl seminar 14462 “Systems and Algorithms for Large-scale Graph Analytics” within a joint session. The organizers are confident that this discussion has led to a better understanding of each other’s community and their contributions. We also hope and think that this exchange will lead to an accelerated dissemination of the respective leading research results across community borders.

Finally, Brugger presented innovative hardware specifically designed to support certain graph algorithms. Talks with a particular focus on innovative applications from outside the core of computer science were presented by several speakers as well. Both Srivastav and Buluc, for example, described algorithms for sequence assembly, a problem in bioinformatics with massive data sets.

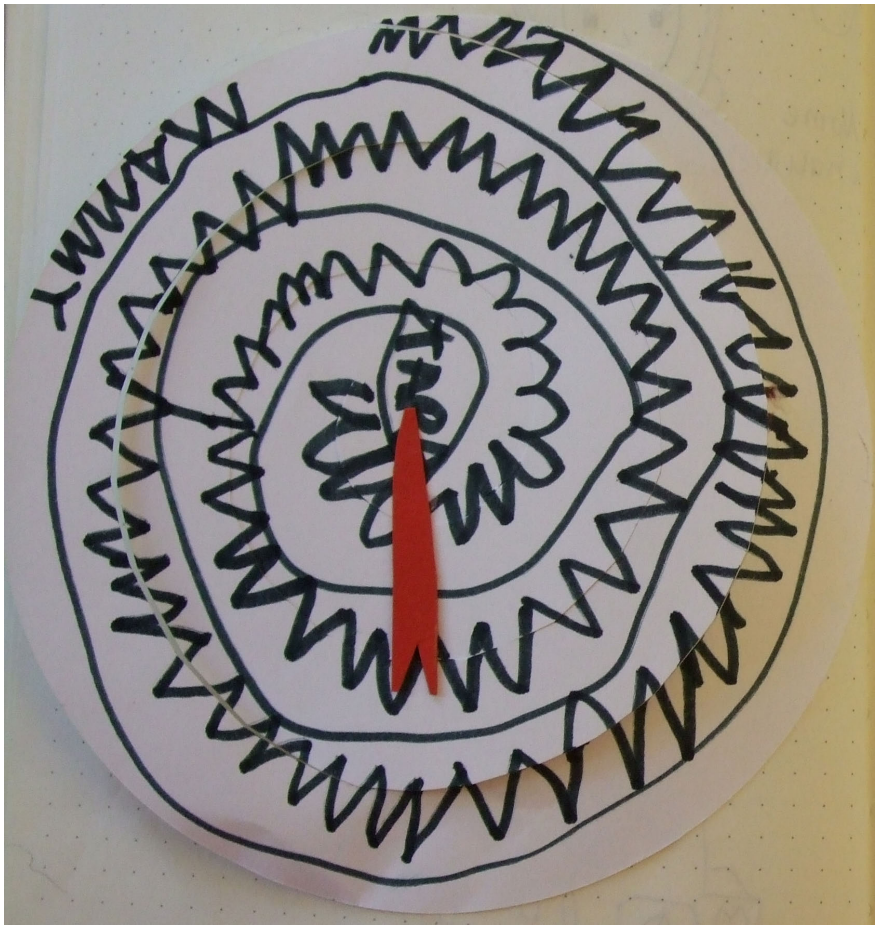

Fig. 4.16
Drawing for the Dagstuhl children's guest book by Tara (5), daughter of Dagstuhl Seminar 14281 participants Kathi Fisler and Shriram Krishnamurthi.

4.67 Systems and Algorithms for Large-scale Graph Analytics

Organizers: Eiko Yoneki, Amitabha Roy, and Derek Murray

Seminar No. 14462

Date: November 9–12, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.59

Participants: Paolo Boldi, Luis Ceze, Hassan Chafi, Lei Chen, Felix Cuadrado, Valentin Dalibard, Khuzaima Daudjee, Šarunas Girdzijauskas, Torsten Grust, Sungpack Hong, Kimberly Keeton, Terence P. Kelly, Arijit Khan, Wolfgang Lehner, Andrew Lenharth, Karthik Nilakant, M. Tamer Özsu, Mark H. Oskin, John Owens, Roger Pearce, Ella Peltonen, Stefan Plantikow, Amitabha Roy, Sebastian Schelter, Alexander Ulrich, Eiko Yoneki, Willy Zwaenepoel

Graph analytics, the class of data analysis that deals with data forming networks, is emerging as a huge consumer of computational resources due to its complex, data-hungry algorithms. Social networking, personal medicine, bioinformatics, text/graphics content analysis and search engines are a few examples where Tera-, Peta- or even Exa-scale graph processing is required. Graph algorithms are becoming increasingly important for solving multiple problems in diverse fields. As these problems grow in scale, parallel computing resources are required to meet the computational and memory requirements. Notably, the algorithms, software and hardware that have worked well for developing mainstream parallel applications are not usually effective for massive-scale graphs from the real world, which exhibits more complex and irregular structures than traditional data in scientific computing.

Research into large scale graph processing within the computer science community is currently at an early and fragmented stage. This seminar brought together researchers from systems, computer architecture, algorithms and databases to discuss emerging trends and to identify opportunities for future advancement. Prior to the seminar, we had prepared a range of research questions below.

1. What is the correct algorithmic abstraction for systems handling large graphs? Algorithmic complexity researchers use PRAM and I/O complexity models to characterize the algorithmic complexity of graph processing. On the other hand, systems researchers largely build systems that implement scatter-gather and label propagation models of computation. These are different world views rendering theory and practice incompatible with each other. We will begin work towards a formal algorithmic model for existing large scale graph processing systems as part of this seminar with a view to answering this question. This model should accurately describe large-scale graph processing systems built by the systems community as well as be formal enough to enable algorithmic complexity researchers to draw useful conclusions about their scalability with data set size. This will require close co-operation between theoretical computer scientists on the algorithms front to talk to practical systems researchers.
2. What is the taxonomy of applications that graph processing systems should support? Can it be reduced to a set of representative benchmarks that researchers in this area need to care about? We can currently identify two main interest areas. The first is large scale graph traversal, of interest to the high performance computing and web-services community; primarily driven by security applications and from data mining needs. The second is spectral approaches, primarily of interest to the machine learning community, building systems such as Graphlab. The output from this agenda item will be a clear set of well defined applications that the community can agree will serve as objects of study for building high performance graph processing systems.
3. What is the correct interface to the system that may be assumed when building a graph DSL? DSL researchers are interested in productive and easy ways to specify graph computations. However they have given relatively little thought to interfacing in an efficient way to systems that execute graph computation. The agenda item therefore will be a discussion between programming language researchers and systems, algorithms and database people researchers about the correct level of interface between a DSL and the underlying systems. A good model in this regard is the decoupling of database systems from ways to query them using declarative languages like SQL. The litmus test for success for this agenda item will therefore be a sketch for a DSL that exposes opportunities for optimization, is productive to use and at the same time is oblivious to the underlying system.

4. What is the design trade-off among different graph processing approaches. For example, (i) the general graph processing system vs. the dedicated approaches specially optimized for specific graph problems, (ii) the running time between pre-processing and graph processing, (iii) performance vs. running expense.

The seminar identified whole graph analytics and point queries on graphs that explore neighborhoods of vertices as distinct application domains, which require separate treatment and systems. All the participants agreed that there was an urgent need to standardize benchmarks and datasets in order to make meaningful progress with graph processing – particularly given the diverse nature of the communities involved. In addition, the seminar identified a number of interesting approaches and trends. There was also considerable participation from industry, which included work in graph databases as well as new systems architectures that will require practitioners to rethink traditional approaches for graph processing.

The seminar consisted of 6 sessions on focused topic presentations and discussions, followed by a joint session with the seminar 14461 on "High-performance Graph Algorithms and applications computational Science". At the last day of the seminar, the whole morning was dedicated to the discussion on the challenges and future directions of large-scale graph processing (see Section Challenges and Future directions).

4.68 Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment

Organizers: Jon Crowcroft, Adam Wolisz, and Arjuna Sathiaseelan

Seminar No. 14471

Date: November 16–21, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.78

Participants: Panayotis Antoniadis, Roger Baig Vinas, Saleem Bhatti, Georg Carle, Weverton Cordeiro, Jon Crowcroft, Roderick Fanou, Michael P. Fourman, Thomas Hühn, Karin Anna Hummel, Renato Lo Cigno, Leonardo Maccari, Mahesh Marina, Leandro Navarro, Irene Ng, Jörg Ott, Fernando M. V. Ramos, Arjuna Sathiaseelan, Henning G. Schulzrinne, Gareth Tyson, Falk Von Bornstaedt, Klaus Wehrle, Michael Welzl, Adam Wolisz, Rüdiger Zarnekow, Marco Zennaro, Anatolij Zubow

Universal Internet Access is considered as one of the fundamental requirements in todays digital age as clean water, roads, schools etc. Enabling universal Internet access is one of the key issues that is being currently addressed at the European level via the Digital Agenda for Europe (DAE) as well as globally. Recognising the importance of broadband Internet, several developed countries have their own national broadband plan, such as the Broadband Development UK (BDUK) in the UK and the National Broadband Plan in the USA.

However a lack of access to the Internet and broadband is a global phenomenon that proportionately and negatively affects the poorest countries in world, where challenges to socio-economic development are most pronounced. It is estimated that only 41% of the worlds households are connected to the Internet. Half of them are in less developed countries, where household Internet penetration has reached 28%. This is in stark contrast to the 78% of households in more developed countries.

The disparity in access is even more worrying when one realises that the positive impact of increased Internet and broadband access is greater than any other ICT. In 2009, the World Bank found that in low and middle-income countries a 10% increase in broadband Internet penetration accelerated economic growth by 1.38%. Moreover, the positive effect of Internet and broadband on economic growth and social development are felt more in less developed countries, like those in sub-Saharan Africa, than in more developed countries, creating opportunities for levelling up and greater equality.

The main barriers to the economic growth and social benefits identified by the World Bank include the cost of services and a lack of access to terrestrial and wireless networks. Indeed, there is general consensus upon the impact of these challenges, especially that of cost. Brahima Sanou, Director of the Telecommunication Development Bureau (BDT) at the ITU notes *Broadband is still too expensive in developing countries, where it costs on average more than 100 per cent of monthly income, compared with 1.5 per cent in developed countries*. There are indeed several challenges (political, regulatory, socio-economical, technological) to the realization of a Future Internet capability that will offer appropriate access to all parts of society.

The goal of our seminar was to bring together an interdisciplinary group of researchers from academia and research organisations as well as industry to understand the different challenges in enabling universal Internet access and to discuss potential solutions for solving some of the challenges.

This report provides an overview of the talks that were given during the seminar. We also had a dedicated breakout session with two groups specifically focussing on *Socio-Economic Models and Role of Community Networks* and *Internet in a box*. We also had longer informal discussions on specific focussed topics. The discussions and outcomes are summarised in this 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 organisation and support.

4

4.69 Multiscale Spatial Computational Systems Biology

Organizers: David Gilbert, Monika Heiner, Koichi Takahashi, and Adelinde M. Uhrmacher

Seminar No. 14481

Date: November 24–28, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.138

Participants: Steven Andrews, Satya Arjunan, Gianfranco Balbo, Ezio Bartocci, Marco Beccuti, Arne Bittig, Mary Ann Blätke, Luca Bortolussi, Lutz Brusch, Alfonso Bueno-Orovio, Kevin Burrage, Francesca Cordero, Walter de Back, Andreas Deutsch, Maciej Dobrzynski, Dirk Drasdo, Radek Erban, Martin Falk, Flavio Fenton, Jérôme Feret, Chiara Fornari, Akira Funahashi, Anna Gambin, David Gilbert, Simon Hardy, Monika Heiner, Mostafa Herajy, Dagmar Iber, Kazunari Kaizu, Andrzej M. Kierzek, Tetsuya J. Kobayashi, Marek Kochanczyk, Fei Liu, Guillaume Madelaine, Wolfgang Marwan, Carsten Maus, Sebastian Nähring, Joachim Niehren, Ovidiu Pârvu, Christian Rohr, Guido Sanguinetti, Koichi Takahashi, Yuichi Togashi, Adelinde M. Uhrmacher

This seminar built on the tradition of two previous Dagstuhl seminars on *Formal Methods in Molecular Biology* in 2009 and 2011 (Seminar 09091, Seminar 11151), but with a special focus on multiscale and spatial modeling and simulation.

Multiscale modelling goes beyond the traditional approach of modelling at just one spatial/temporal scale or organizational level. Until now most models have largely ignored locality within the cell, or cell-cell interactions. However, with the insight that spatial phenomena like localisation and crowding have a considerable influence on reaction processes and many processes cannot be understood with reference to one organisation level only (intra- or inter-cellular dynamics), the need for effective and efficient modelling and simulation approaches arises.

The challenges for computer science and mathematics include the development of suitable modelling approaches and associated tools to create coherent descriptions of biological systems by integrating several spatial and/or temporal scales, and methods for the simulation and analysis of the models.

The overall motivation for this seminar was the exploration of the most recent advances in these methods. The seminar brought together researchers working in modelling and analysis of biological systems with diverse professional backgrounds, including informaticians, mathematicians, engineers, biologists, physicians.

A distinguishing factor of the seminar was the modelling exercise – where teams explored different modelling paradigms, in order to better understand the details of the approaches, their challenges, potential applications, and their pros and cons. This activity was carried out in a collaborative and self-directed manner using the Open Space Technology approach as evidenced by a high degree of communication both within and between the teams. Eight teams were formed, and reports from five of them are included in full report. The teams were formed around the following focii:

- Small GTP-ase pathway.
- Continuous multiscale models for biological tissue.
- Simulating macromolecular crowding with particle and lattice-based methods
- Multiscale modeling of S1P metabolism, secretion and signaling
- DNA structural dynamics.
- Dictyostelium discoideum: Aggregation and Synchronisation of Amoebas in Time and Space.
- Towards a standard exchange format for spatial, multilevel multicellular models.
- Model checking for multiscale spatial biological systems.

The participants decided to take forward the activities in the future outside Dagstuhl, with the goals of carrying out collaborative research, producing scientific papers and applying for larger scale funded international research projects.

4.70 Automated Planning and Model Checking

Organizers: Alessandro Cimatti, Stefan Edelkamp, Maria Fox, Daniele Magazzeni, and Erion Plaku

Seminar No. 14482

Date: November 23–28, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.11.227

Participants: Steven Andrews, Sergiy Bogomolov, Dragan Bosnacki, Ronen I. Brafman, Alessandro Cimatti, Stefan Edelkamp, Georgios Fainekos, Maria Fox, Malte Helmert, David Hsu, Alan Hu, Hanna Kurniawati, Derek Long, Daniele Magazzeni, Sheila McIlraith, Bernhard Nebel, Doron A. Peled, Erion Plaku, Andreas Podelski, Franco Raimondi, Sylvie Thiébaux, Martin Wehrle, Anton Wijs, Brian C. Williams, Paolo Zuliani

In the area of formal methods, model checking deals with the problem of fully-automated property validation and correctness verification. Given a formal model of a system and a property specification, the task is to explore the state space and verify whether or not the property is satisfied by the model. In artificial intelligence, automated planning deals with the automatic generation of plans for achieving a goal. Given the description of the initial state, the goal state, and the set of possible actions, a planner uses heuristic search to look for a sequence of actions that transforms the initial state into the goal state.

There has been a lot of work on the exchanges between the two areas (automated planning and model checking), based on the observation that a model-checking problem can be cast as a planning problem, where the goal state is a state violating the property to be verified in the model-checking problem. Thus, if a plan is found by the planner, it corresponds to error trace that a model checker would return (this paradigm is called directed model checking). The link can also be exploited in the other way around, using a model checker to search the planning state space, stopping the search when a goal state is found (this paradigm is called planning via model checking).

The general aim of this Dagstuhl Seminar was to increase the synergy between the two research communities. This involved sharing views, thoughts and contributions across the following streams:

Techniques and Tools: During the seminar we considered the most recent advances in automated planning and model checking and explored the possibility of using recent planning tools (heuristic search, sampling-based motion-planning algorithms, symbolic search algorithms, ...) for system falsification (particularly in challenging domains such as hybrid systems) and for boosting the use of model checking systems for finding plans.

Modelling Languages: One of the goals of the seminar was to consider the family of PDDL languages and the formalisms for describing verification problems and discuss how to make the communication between the two areas easier, exploring the possibility of common languages or translation between existing formalisms.

There were a number of talks on "X" Modulo Theories, where "X" ranged from SAT and Search to Planning. These talks explored the relationship between generic solution methods and different proof systems, leading to discussions about the relative benefits of viewing problems from the perspectives of the different generic methods. This fostered an improved understanding of each other's perspectives and modelling approaches.

Discussions also focused on relationships between hybrid planning and hybrid model-checking. There were tutorials on the modelling languages used in the two paradigms, and their semantics, and the discrepancies between these led to lively discussion. In hybrid planning using PDDL, a key semantic issue is the use of epsilon time to separate inter-dependent actions, in order to prevent the planner from relying on synchronised activity. For example, if an action, A, achieves the precondition of another action, B, the validity of the plan depends on A being ordered strictly before B, by at least epsilon time. This is because the state following the co-occurrence of these two actions is indeterminate. The model-checking community does not require this epsilon. In hybrid model-checking a partial order on events is maintained, and there was an extended discussion about why planning forces an ordering using epsilon separation when this is not necessary in model-checking. From this discussion the following distinction emerged: model-checking simply requires there to be a single ordering of events that is consistent with the constraints, as this provides the required counter-example to the correctness of the model. In planning, by contrast, *all* orderings of events must be consistent with the constraints, requiring exponential work to

check the validity of a partially ordered plan. Once this point was understood there was a greater common understanding between the planning and model-checking proponents, and greater appreciation of the crossover between modelling languages and methods.

Other topics covered by the contributed talks include:

- directed model checking and falsification
- plan validation
- heap and other data structures
- GPU-based state space exploration
- hybrid systems
- heuristic search
- planning and verification on real-world scenarios

The program featured the following components:

- On Monday we started with 7 tutorial-type introductory talks about plans validation, planning in hybrid systems through model checking, guided search for hybrid systems, falsification of hybrid systems through motion planning, planning via symbolic model checking, directed model checking of timed systems, heap implementations. The purpose of these tutorials was to familiarise members of the different communities with the basics of the other fields and with the existing synergies between the fields.
- From Tuesday, each day featured one or two long talks plus a number of short talks, with enough time for discussion after each talk.
- On Tuesday and Thursday afternoon we had two open discussion sessions.
- Wednesday afternoon featured a hike.

A feature of this seminar was the very high level of engagement and interaction between the participants, leading to a lively and productive week. The decision not to formalise discussions into panels or break-out sessions proved to be a good one, allowing more flexible response to topics as they arose. Similarly, the decision to leave some of the talk slots open allowed spontaneous pursuit of ideas that came out of discussions. The mixture of long and short talks also encouraged this. The workshop ended on a high note, with many new ideas for collaboration having been identified.

4.71 Socio-Technical Security Metrics

Organizers: Dieter Gollmann, Cormac Herley, Vincent Koenig, Wolter Pieters, and Martina Angela Sasse

Seminar No. 14491

Date: November 30 to December 5, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.12.1

Participants: Zinaida Benenson, Sören Bleikertz, Rainer Böhme, Tristan Caulfield, Kas P. Clark, Trajce Dimkov, Simon N. Foley, Carrie Gates, Dieter Gollmann, Dina Hadziosmanovic, Carlo Harpes, Cormac Herley, Roeland Kegel, Vincent Koenig, Stewart Kowalski, Aleksandr Lenin, Gabriele Lenzini, Mass Soldal Lund, Sjouke Mauw, Daniela Oliveira, Frank Pallas, Sebastian Pape, Simon Parkin, Shari Lawrence Pfleeger, Wolter Pieters, Kai Rannenberg, Roland Rieke, Martina Angela Sasse, Paul Smith, Ketil Stolen, Axel Tanner, Sven Übelacker, Michel van Eeten, Jan Willemson, Jeff Yan

Introduction

Socio-technical vulnerabilities

Information security, or cyber security, is not a digital problem only. Humans have been termed "the weakest link", but also physical access plays a role. Recent cyber attacks cleverly exploit multiple vulnerabilities of very different nature in the socio-technical systems that they target. For example, the StuxNet attack relied both on Industrial Control System (ICS) vulnerabilities and on the physical distribution of infected USB sticks, allowed by the business processes in the target facilities [8]. With new developments such as cloud computing, the attack surface of the systems only increases, and so do the options for potential attackers. At any company in the service supply chain, there may be malicious insiders or benevolent employees who fall victim to social engineering, and they influence the security of the system as a whole significantly. In order to compare and prioritize attacks and countermeasures, for example in terms of risk, the different types of vulnerabilities and threats need to be expressed in the same language. The seminar on "Socio-technical security metrics" aims at developing cross-domain metrics for this purpose.

Defining metrics

The idea of defining information security in terms of risk already appeared quite a while ago [2, 10]. Since then, many metrics have been proposed that aim to define attacks and attack opportunities in information systems in quantitative terms (see e.g. [7, 12]). Often, likelihood and impact of loss are mentioned as the key variables, from which risk can then be calculated. Furthermore, notions of vulnerability, difficulty, effort, cost, risk for the attacker, and many more, show up in the literature.

Even in a purely technical setting it is not always clear how all these different concepts are related. Still, including the human element forms a particular challenge, which deserves a separate event and a better integrated community. Too often it is thought that models of humans in the social sciences and models of technology are fundamentally incompatible. This inhibits progress on some very relevant questions: How does sending a phishing message compare to an SQL injection, in terms of the above mentioned variables? Or do we need additional notions in the technical models to express the human elements, or in the social science models to express the technical ones?

We thus need unified – or at least comparable – metrics that apply to all types of vulnerabilities. In order to represent socio-technical attacks, the key concepts need to apply to very different types of actions in an attack, including technical exploits and social engineering alike. This requires knowledge on technical infrastructures, social science, and actual incidents. To enable meaningful socio-technical security metrics, key features to be addressed in the seminar are outlined below.

Multi-step attacks

Cyber attacks, like StuxNet, tend to consist of multiple steps, combining technical and social or organizational vulnerabilities. Attack trees [17] are often used to represent possible multi-step attacks on systems, and they can be annotated with quantitative metrics. It has also been proposed to develop formal analysis techniques and simulations ("attack navigators") that generate such trees based on a model of the socio-technical system at hand [5, 16]. By defining methods to calculate metrics for attacks from metrics for steps, one can compare the attacks in terms of the metrics, e.g. difficulty. However, next to methods for prediction, one would also want to be able to estimate the relevant parameters for the model based on observed events. For example, if one observes a set of successful and unsuccessful attacks, what does that say about the difficulty of the steps involved, and how does that influence the prediction of possible future events? Statistical methods from social science may assist here [15].

Estimating metrics from data

Data is thus key to developing good metrics, but obtaining them requires care. Given the data that is typically available in organizations already, including enterprise architecture, network logs, and potentially even organizational culture, how to obtain the right metrics from that data? What could be the role of "Big Data" in improving security metrics? And how to acquire additional data in tailor-made experiments? From the modeling point of view, a distinction can be made here between bottom-up approaches, leveraging existing data, and top-down approaches, defining targeted data collection methods and experiments. A good example on the social side are the phishing studies by Jakobsson & Ratkiewicz [6]. On the technical side, intrusion detection systems may constitute an important source of data.

Attacker models

As security threats originate from attackers and not from nature, attacker models are key for security metrics [9]. Attackers will adapt their strategies to the security situation, and also to newly deployed countermeasures. We therefore need meaningful and measurable features of attackers that can be used as a basis for the metrics. For example, the motivation of an attacker may determine the goal of the attack, the resources available to an attacker may determine the number of attacks that he can attempt, and attacker skill may determine the likelihood of success. Costs of an attack as well as risk of detection influence attacker behavior [3]. Again, the theoretical and empirical basis of such models needs to be carefully studied, and (security) economics may provide important insights here.

Countermeasures

All these aspects come together in one final goal: supporting investments. In order to estimate the cost-effectiveness of security measures (also called ROSI, for return on security investment), one would need metrics for both the risk prevented by the countermeasures, and of their cost. The former could be calculated based on the properties discussed above. The latter, however, is far from trivial by itself, as costs not only involve investment, but also operational costs. Operational costs, in turn, may include maintenance and the like, but an important factor in the total cost of ownership is impact on productivity. Security features may increase the time required to execute certain tasks, and people have a limited capacity for complying with security policies. If security is too cumbersome or misdirected, people will find workarounds, and this may reduce the effect of the measures on risk [1]. Thus, metrics for countermeasure cost form an important topic in itself, requiring input from the human factors and usable security domains.

Another application area for the metrics would be selection among alternative system designs. For example, if two vendors offer the same equipment or service, but one is much cheaper, how to take security risk into account when making this decision? Both vendors as well as customers may be interested in security metrics from this point of view. However, metrics would need to be designed carefully in order to avoid creating perverse incentives, tweaking systems to score high on the metrics without actually being "better".

Communities

In order to develop meaningful metrics for socio-technical security, participants from the following communities were invited:

- Security metrics and data-driven security, for obvious reasons;
- Security risk management, to provide input on suitable risk variables to be included;
- Security economics, to build upon economic theories of behavior of both attackers and defenders;
- Security architectures, to get relevant data on information system architecture and incidents;
- Formal methods, to analyze attack opportunities in complex systems;
- Social / crime science, to understand attacker behavior and the influence of controls;
- Human factors, to understand the impact of security controls on users.

Main findings

Paraphrasing some ancient philosophical questions (what is there, what can we know, what should we do), we can structure the main outcomes of this seminar as follows:

1. What properties are we interested in?
2. What can we measure?
3. What should we do with the measurements?

What properties

One of the main outcomes of the seminar is a much better view on which types of security metrics there are and for which purposes they can be used.

This leads to a distinction between metrics that exclude the real-life threat environment (type I) and metrics that include the real-life threat environment (type II). Metrics describing difficulty or resistance are typically of type I. They give a security metric that is independent of the actual activity of adversaries, or of the targets that they might be after. For example, which percentage of the people fall for a simulated phishing mail. This is similar to what Böhme calls "security level" [4]. The threat environment is often specified explicitly in such metrics, and the metrics may thus enumerate threat types. However, they do not estimate their occurrence rates, and in fact the occurrence rate is often controlled. In the phishing case, the researchers control the properties and occurrence of the phishing e-mails, and describe the e-mail (controlled threat) in their results.

Metrics describing loss (risk) or incidents are typically of type II. They describe undesired events that happen based on interaction of the system with a threat environment (activity of adversaries), and their consequences. For example, the number of infected computers of a particular Internet Service Provider [18].

An illustration of this difference is the following. Consider two systems, system A and system B [13]. In system A, a locked door protects € 1,000. In system B, an identical locked door protects € 1,000,000. Which system is more secure? Or, alternatively, which door is more secure? One might say that system A is more secure, as it is less likely to be attacked (assuming the attacker knows the system). On the other hand, one might say that the doors are equally secure, as it is equally difficult to break the lock. The former argument is based on including an evaluation of the threat environment, the latter on excluding it.

Obviously, when trying to derive type II metrics from type I metrics, one needs metrics on the threat environment as well. For example, when one wants to calculate risk related to phishing attempts, and one knows how likely one's employees are to fall for phishing mails based on their sophistication, then one also needs information on the expected frequency of phishing mails of certain levels of sophistication in order to calculate the risk. Such models of the threat environment may be probabilistic or strategic (game-theoretic), representing non-adaptive and adap-

tive attackers, respectively. Probabilistic models, in turn, may be either frequentist (based on known average frequencies) or Bayesian (based on subjective probabilities). The various points of view have not been fully reconciled up to this point, although integration attempts have been made [14].

Another consideration is the integration of security metrics from different domains: digital, physical and social. Often, there are different styles of type I metrics, which one would like to integrate in a single type II metric representing the level of risk in a socio-technical system (e.g. an organization). Digital metrics may represent difficulty as required skill (e.g. CVSS), physical metrics may use required time (e.g. burglar resistance), and social metrics may use likelihood of success (e.g. likelihood of success of phishing attempts). Integration of these metrics is still an open challenge.

What measurements

The seminar discussed methods applied in different scientific communities for measurement purposes. Some of those methods rely on quantitative indicators, some rely on qualitative indicators, and some combine both. A further distinction can be made between subjective and empirical metrics, e.g. expert judgements versus monitoring data. Hereafter, and for the purpose of illustration, we have drawn a non-comprehensive list of such methods. They can be applied individually or in a complementary way, covering one measure or combined measures. A specific usage we consider underrepresented so far is the combination of methods in an effort to augment the measurement quality, or to provide information about the validity of a new measure. This approach has often been referred to, during the seminar, as triangulation of measures.

These are social methods discussed in the seminar:

- semi-structured interviews; in-depth interviews; surveys;
- observations of behavior;
- critical incident analysis;
- laboratory experiments; field experiments;
- expert / heuristic analysis / cognitive walkthrough;
- root cause analysis.

These are technical methods discussed in the seminar:

- security spending;
- implemented controls;
- maturity models;
- incident counts;
- national security level reports;
- service level agreements.

It is important to assess which type of metric (type I or type II) is produced by each of the techniques. For example, penetration testing experiments produce type I metrics, whereas incident counts produce type II. Maturity models and national security level reports may be based on a combination of type I and type II metrics. In such cases, it is important to understand what the influence of the threat environment on the metrics is, in order to decide how the metrics can be used.

What usage

Security metrics can contribute to answering questions about a concrete system or questions about a design (hypothetical system), and questions about knowledge versus questions about preferences. Here, we focus on a simpler distinction, namely between knowledge and design questions. In the case of knowledge questions, metrics are used to gather information about the world. In the case of design questions, metrics are used to investigate a design problem or to evaluate the performance of a design, such as a security control. In terms of knowledge questions, a typical usage discussed is a better understanding of the human factor in security. In terms of design, possible questions are how much security feedback a system should give to users or operators, or how to provide decision support for security investment.

Security metrics may have several limitations. In particular, many metrics suffer from various forms of uncertainty. It may be unclear whether the metrics measure the right thing (validity). Even if this is the case, random variations may induce uncertainty in the values produced (reliability). It is therefore important to understand the implications of such uncertainties for decisions that are made based on the metrics. Triangulation may contribute to the reduction of uncertainty. In some cases, quantitative metrics may not be possible at all, and qualitative methods are more appropriate.

Another limitation is that stakeholders may behave strategically based on what they know about the metrics (gaming the metrics). If stakeholders are rewarded when their security metrics become higher, they may put effort into increasing the metrics, but not “actual security”. Even if the metrics are valid under normal circumstances, this needs not be the case under strategic behavior.

Conclusions

Security is difficult to measure, which should not be a surprise to those involved. However, to understand security in today’s complex socio-technical systems, and to provide decision support to those who can influence security, rigorous conceptualisation, well-defined data sources and clear instructions for use of the metrics are key assets. This seminar laid the foundations for understanding and applying socio-technical security metrics.

In particular, we strove for clarity on (a) the different types of security metrics and their (in)compatibility, (b) the different sources and methods for data extraction, and (c) the different purposes of using the metrics, and the link with types, methods and sources. Several papers are planned as follow-up activities, as described in the reports of the working groups (see full report). On many topics there are different views, which may not always be compatible, as was clear from the panel discussion. Future follow-up seminars would be very valuable to address the open problems which are described in the full report.

References

1 A. Beautement, M. A. Sasse, and M. Wonham. The compliance budget: Managing security behaviour in organisations. In *Proc. of the 2008 Workshop on New Security Paradigms*, NSPW’08, pp. 47–58, New York, NY, USA, 2008. ACM.

2 B. Blakley, E. McDermott, and D. Geer. Information security is information risk management. In *Proc. of the 2001 New Security Paradigms Workshop*, pp. 97–104, New York, NY, USA, 2001. ACM.

3 A. Buldas, P. Laud, J. Priisalu, M. Saarepera, and J. Willemson. Rational choice of security measures

via multi-parameter attack trees. In *Critical Information Infrastructures Security*, volume 4347 of *LNCS*, pp. 235–248. Springer, 2006.

4 R. Böhme. Security metrics and security investment models. In Isao Echizen, Noboru Kunihiro, and Ryoichi Sasaki, editors, *Advances in Information and Computer Security*, volume 6434 of *LNCS*, pp. 10–24. Springer, 2010.

5 T. Dimkov, W. Pieters, and P. H. Hartel. Portunes: representing attack scenarios spanning through the physical, digital and social domain. In *Proc. of the Joint Workshop on Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security (ARSPA/WITS'10)*, volume 6186 of *LNCS*, pp. 112–129. Springer, 2010.

6 P. Finn and M. Jakobsson. Designing ethical phishing experiments. *Technology and Society Magazine, IEEE*, 26(1):46–58, 2007.

7 M. E. Johnson, E. Goetz, and S. L. Pfleeger. Security through information risk management. *IEEE Security & Privacy*, 7(3):45–52, May 2009.

8 R. Langner. Stuxnet: Dissecting a cyberwarfare weapon. *Security & Privacy, IEEE*, 9(3):49–51, 2011.

9 E. LeMay, M. D. Ford, K. Keefe, W. H. Sanders, and C. Muehrcke. Model-based security metrics using adversary view security evaluation (ADVISE). In *Proc. of the 8th Int'l Conf. on Quantitative Evaluation of Systems (QEST'11)*, pp. 191–200, 2011.

10 B. Littlewood, S. Brocklehurst, N. Fenton, P. Mellor, S. Page, D. Wright, J. Dobson, J. McDermid, and D. Gollmann. Towards operational measures of computer security. *Journal of Computer Security*, 2(2–3):211–229, 1993.

11 H. Molotch. *Against security: How we go wrong at airports, subways, and other sites of ambiguous danger*. Princeton University Press, 2014.

12 S. L. Pfleeger. Security measurement steps, missteps, and next steps. *IEEE Security & Privacy*, 10(4):5–9, 2012.

13 W. Pieters. Defining "the weakest link": Comparative security in complex systems of systems. In *Proc. of the 5th IEEE Int'l Conf. on Cloud Computing Technology and Science (CloudCom'13)*, volume 2, pp. 39–44, Dec 2013.

14 W. Pieters and M. Davarynejad. Calculating adversarial risk from attack trees: Control strength and probabilistic attackers. In *Proc. of the 3rd Int'l Workshop on Quantitative Aspects in Security Assurance (QASA)*, LNCS, Springer, 2014.

15 W. Pieters, S. H. G. Van der Ven, and C. W. Probst. A move in the security measurement stalemate: Elo-style ratings to quantify vulnerability. In *Proc. of the 2012 New Security Paradigms Workshop*, NSPW'12, pages 1–14. ACM, 2012.

16 C. W. Probst and R. R. Hansen. An extensible analysable system model. *Information security technical report*, 13(4):235–246, 2008.

17 B. Schneier. Attack trees: Modeling security threats. *Dr. Dobb's journal*, 24(12):21–29, 1999.

18 M. J. G. Van Eeten, J. Bauer, H. Asghari, and S. Tabatabaie. The role of internet service providers in botnet mitigation: An empirical analysis based on spam data. OECD STI Working Paper 2010/5, Paris: OECD, 2010.

4.72 The Synergy Between Programming Languages and Cryptography

Organizers: Gilles Barthe, Michael Hicks, Florian Kerschbaum, and Dominique Unruh

Seminar No. 14492

Date: November 30 to December 5, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.12.29

Participants: Joseph Ayo Akinyele, David Archer, Manuel Barbosa, Gilles Barthe, Karthikeyan Bhargavan, Bruno Blanchet, Dan Bogdanov, Niklas Büscher, Stephen Chong, Véronique Cortier, Francois Dupressoir, Cédric Fournet, Matthew Hammer, Michael Hicks, Catalin Hritcu, Stefan Katzenbeisser, Florian Kerschbaum, Boris Köpf, Markulf Kohlweiss, Sven Laur, Alex Malozemoff, Sarah Meiklejohn, Esfandiar Mohammadi, Axel Schröpfer, Alley Stoughton, Eran Tromer, Dominique Unruh, Santiago Zanella-Béguelin

The seminar schedule consisted of three components: short, two minute introduction talks (one for each participant), longer technical talks and open discussions on four different subjects. The first two days consisted of the introduction talks, followed by most of the technical talks. The seminar attendees had a mix of backgrounds, with one half (roughly) leaning heavily toward the PL (programming languages) side, and the other half leaning more towards the crypto side. The diversity of talks reflected this diversity of backgrounds, but there was much opportunity to meet in the middle and discuss open problems. The latter days mixed some remaining technical talks with open discussion sessions focusing on various problems and topics.[25] In particular, participants voted to select four breakout sessions: Secure Computation Compilers, Crypto verification, Obfuscation, and Verified implementations.

This section summarizes some interesting discussions from the seminar, in three parts. First, we consider the activities involved in developing programming languages the interface with cryptography, and surveying the research of the seminar participants. Second, we explore how reasoning in PL and Crypto compare and contrast, and how ideas from one area might be relevant to the other. Finally, we survey open problems identified during the discussions.

Programming languages for cryptography

One connection emerged repeatedly in the talks and discussions: the use of programming languages to do cryptography, e.g., to implement it, optimize it, and prove it correct.

Programming languages can be compiled to cryptographic mechanisms Programming languages can make cryptographic mechanisms easier to use. For example, the systems Sharemind, ShareMonad, CBMC-GC, and Wysteria are all designed to make it easier for programmers to write secure multiparty computations (SMCs).

In an SMC, we have two (or more) parties X and Y whose goal is to compute a function F of their inputs x and y, whereby each party only learns the output $F(x, y)$, but does not "see" the inputs. Cryptographers have developed ways to compute such functions, such as garbled circuits[26] and computing on secret shares[27], without need of a trusted third party. These systems shield the programmer from the workings of these mechanisms, compiling normal-looking programs to use the cryptography automatically. The languages can also provide additional benefits, such compiler-driven optimization.

This line of work is motivated by privacy- and/or integrity-preserving outsourcing of computation, e.g., as promised by The Cloud. Programming languages have been designed to compile to

[25] As a break from the technical program, we went on a group outing to Trier on Wednesday afternoon, where we enjoyed a guided historical tour and enjoyed the city's Christmas market.
[26] https://www.usenix.org/legacy/event/sec11/tech/full_papers/Huang.pdf
[27] http://www.math.ias.edu/~avi/PUBLICATIONS/MYPAPERS/GMW87/GMW87.pdf
[28] https://eprint.iacr.org/2014/976.pdf
[29] http://eprint.iacr.org/2013/507
[30] http://amiller.github.io/lambda-auth/

other kinds of crypto aside from SMC, like zero-knowledge proofs and authenticated data structures. Examples include Geppetto[28], SNARKs for C[29] and LambdaAuth[30].

Combinations also exist, such as compiling to support Authenticated SNARKs.

Programming languages for implementing cryptography. The above languages aim to make computations secure through the use of cryptography, introduced by the language's compiler. We are also interested in implementing the cryptographic algorithms themselves (e.g., for symmetric or public key encryption). The implementation task could be made easier, more efficient, or more secure by employing a special-purpose language. Two representatives in this space are CAO[31] and Cryptol[32]. Both are domain-specific, and both make it easier to connect implementations to tools for automated reasoning. The Seminar also featured work on synthesizing cryptography (block ciphers) from constraint-based specifications.[33]

Programming languages methods to prove security of cryptographic protocols and/or their implementations. When a cryptographer defines a cryptographic protocol, she must prove it is secure. Programming languages methods can be used mechanically confirm that a proof of security is correct. Systems like ProVerif[34], CryptoVerif[35], EasyCrypt[36] and CertiCrypt[37] support cryptographic protocol verification, with varying kinds of assurance. These systems build on ideas developed in general verification systems like Coq or Isabelle.

Likewise, when a programmer implements some cryptography (in a language like C), she would like to formally verify that the implementation is correct (no more Heartbleed!). For example, we'd like to know that an implementation does not have side channels, it uses randomness sufficiently, it has no buffer overflows, etc. Once again, verification can be achieved using tools that are underpinned by PL methods developed in formal verification research. Frama-C[38] and Fstar[39] have been used to verify implementations.

Formal reasoning for PL and Crypto.

Beyond using PLs as a tool for easier/safer use of Crypto, there is an opportunity for certain kinds of thinking, or *reasoning*, to cross over fruitfully between the PL an Crypto communities. In particular, both communities are interested in formalizing systems and proving properties about them but they often use different methods, either due to cultural differences, or because the properties and systems of interest are simply different. During the seminar we identified both analogous, similar styles of reasoning in two communities and connection points between the different styles of reasoning.

Analogies between PL and Crypto reasoning. The Ideal/Real paradigm was first proposed by Goldreich, Micali, and Widgerson in their work on Secure Multiparty Computation (SMC) [3, 4], and further developed by Canetti in his universal composability (UC) framework[40]. The basic idea is to treat a cryptographic computation among parties as if it were being carried out by a trusted third party (the "ideal"), and then prove that the actual implementation (the "real") emulates this ideal, in that the parties can learn nothing more than they would in a protocol involving a trusted party. (The paradigm also handles correctness, robustness, and other emergent properties.)

This is a classic kind of abstraction also present in formal verification: If a program P uses a module M that implements specification S, then relinking P to use M', which also implements S, should preserve the correct execution of P. One talk, by Alley Stoughton, made the interesting observation that the Real/Ideal notion might be a suitable organizing principle around which to verify software is secure, essentially by using the Ideal as a richer kind of security property than is typical in PL (which often looks at properties like information flow control), and using abstraction in key ways to show it is enforced.

In the Crypto setting, the Real-to-Ideal connection is established probabilistically, considering a diminishing likelihood that a computationally bounded adversary would be able to tell the difference between the Real and Ideal. In the PL setting, the specification-implementation connection is established using methods of formal reasoning and logic, and usually without considering an adversary.

However, a notion of adversary does arise in PL-style reasoning. In particular, an adversary can be expressed as a context $C[\cdot]$ into which we place a computation e of interest that is subject to that adversary; the composition of the two is written $C[e]$. One PL property in this setup with a Crypto connection is contextual equivalence, which states that e and e' are equivalent iff for all contexts C the outcome of running $C[e]$ is the same as running $C[e']$ – e.g., both diverge or evaluate to the same result. In a PL setting this property is often of interest when proving that two different implementations of the same abstract data type have the same semantics (in all contexts). In a security setting we can view the contexts as adversaries, and e and e' as the Real and Ideal.

Another useful property is *full abstraction*.[41] This property was originally introduced to connect an operational semantics to a denotational semantics – the former defines a kind of abstract machine that explains how programs compute, while the latter denotes the meaning of a program directly, in terms of another mathematical formalism (like complete partial orders). Both styles of semantics have different strengths, and full abstraction connects them: it requires that e and e' are observationally equivalent (according to the operational semantics) if an only if they have the same denotation (according to the denotational semantics).

In a Crypto setting, we might view the operational semantics as the Ideal and the denotational semantics as the Real, and full abstraction then states that despite the added observational power of the Real setting, an adversary cannot distinguish any more

[31] http://haslab.uminho.pt/mbb/software/cao-domain-specific-language-cryptography
[32] https://galois.com/project/cryptol/
[33] https://eprint.iacr.org/2014/774
[34] http://prosecco.gforge.inria.fr/personal/bblanche/proverif/
[35] http://prosecco.gforge.inria.fr/personal/bblanche/cryptoverif/
[36] https://www.easycrypt.info/trac/
[37] http://certicrypt.gforge.inria.fr/
[38] http://frama-c.com/
[39] http://research.microsoft.com/en-us/projects/fstar/
[40] https://eprint.iacr.org/2000/067.pdf
[41] http://users.soe.ucsc.edu/~abadi/Papers/paper-csf-long.pdf

programs (i.e., learn any additional information) than he could in the Ideal setting. As a recent example of its use, Abadi and Plotkin used full abstraction to reason about the effectiveness of address space randomization. Another recent result is a fully abstract compiler from a type-safe high-level language to Javascript[42]; the compiler effectively defines the denotational semantics, and the fact that it is fully abstract means that the added adversarial power that Javascript provides cannot violate the source language's semantics.

Connections between PL and Crypto. The seminar also brought out ways that PL-style reasoning can be connected to Crypto-style reasoning for stronger end-to-end assurance of security. One connection point was at the Real/Ideal boundary. In particular, for privacy-preserving computation (or computation preserving some other security property), Crypto-style reasoning can first be used to establish that the Real emulates the Ideal, and then PL-style reasoning can consider the security of the Ideal itself.

For example, consider the setting of SMC. Here, we have two (or more) parties X and Y that wish to compute a function F of their inputs x and y, whereby each party only learns the output $F(x, y)$, but does not "see" the inputs. That is, the security goal is to establish that the Real computation of $F(x, y)$ is indistinguishable from the Ideal model of executing F at a trusted third party. While Crypto can establish that a technique like garbled circuits effectively emulates a trusted third party, it does not establish that the output of F, even when computed by the Ideal, does not reveal too much information. For example, if $F(x, y) = y$ then X learns Y's value y directly. More subtly, if $F(x, y) = x > y$, then if $x = 1$, an output of TRUE tells X that Y's value $y = 0$. PL-style reasoning can be applied to functions F to establish whether they are sufficiently private, e.g., by using ideas like knowledge-based reasoning[43] or type systems for differential privacy.[44] PL-style reasoning about knowledge can also be used to optimize SMCs by identifying places where a transformation would not affect security (e.g., no more is learned by an adversary observing the transformed program), but could improve performance.[45]

Another way to connect PL to Crypto is to factor security-sensitive computations into general-purpose and cryptographic parts. Then PL-style methods can be used to specify the overall computation with the Crypto parts carefully abstracted out. The proof of security then follows a PL approach, assuming guarantees provided by the Crypto parts, which are separately proved using Crypto techniques. In a sense we can think of the PL techniques as employing syntactic/symbolic reasoning, and the Crypto techniques employing computational/probabilistic reasoning.

This is the approach taken in LambdaAuth, a language extension for programming *authenticated data structures* (in the style of Merkle trees), in which the key idea involving the use of cryptographic hashes was abstracted into a language feature, and the proof of security combined a standard PL soundness proof along with a proof of the assumption that hash collisions are computationally difficult to produce. Recent work by Chong and Tromer on proof-carrying data similarly considers a language-level problem and proves useful guarantees by appealing to abstracted cryptographic mechanisms.[46] Likewise, work on Memory Trace Obliviousness reasons about Oblivious RAM abstractly/symbolically in a PL setting to prove that the address trace of a particular program leaks no information.[47]

Open problems

Beyond work that is being done, one goal of the seminar was to identify possible collaborations on future work. PL researchers and cryptographers work on common problems from different points of view, so one obvious next step is to collaborate on these problems.

One relevant problem is *side channels*. Cryptographers are concerned with side channels in their implementations, e.g., to make sure the time, space, or power consumption during an encryption/decryption operation does not reveal anything about the key. Likewise, PL folk care about side channels expressed at the language level, e.g. work by Andrew Myers' group on timing channels[48]. Both groups bring a useful perspective.

Another common problem is *code obfuscation*. It was cryptographers that proved that virtual black box (VBB) obfuscation is impossible[49], and proposed an alternative indistinguishability-based definition. PL researchers, on the other hand, have looked at language-oriented views of obfuscation effectiveness, e.g., based on abstract interpretation[50]. Just as the halting problem is undecidable, but practical tools exist that prove termination.[51] I believe that there is an opportunity here to find something useful, if not perfect.

Finally, the question of *composability* comes up in both Crypto and PL: Can we take two modules that provide certain guarantees and compose them to create a larger system while still ensuring properties proved about each module individually? Each community has notions for composability that are slightly different, though analogous, as discussed above. Can we make precise connections so as to bring over results from one community to the other? Crypto currencies, exemplified by BitCoin, are an area of exploding interest. An interesting feature about these currencies is that they provide a foundation for fair, secure multiparty computation, as demonstrated by Andrychowicz, Dziembowski, Malinowski, and Mazurek in their best paper at IEEE Security and Privacy 2014 [1, 2]. Could PL-style reasoning be applied to strengthen the guarantees provided by such computations? Cryptographic properties are often proved by making probabilistic statements about a system subject to a computationally bounded adversary. Could program analyses be designed to give probabilistic guarantees, drawing on the connection between adversary and context mentioned above, to thus speak more quantitatively about the chances that a property is true, or not, given the judgment of an analysis? How might random testing,

[42] http://research.microsoft.com/en-us/um/people/nswamy/supp/full-abstraction.html
[43] http://www.cs.umd.edu/~mwh/papers/mardziel12smc.html
[44] http://www.cis.upenn.edu/~ahae/papers/dfuzz-popl2013.pdf
[45] http://www.cs.umd.edu/~mwh/papers/rastogi13knowledge.html
[46] https://eprint.iacr.org/2013/513
[47] http://www.cs.umd.edu/~mwh/papers/liu13oblivious.html
[48] http://www.cs.cornell.edu/andru/papers/pltiming.html
[49] https://www.iacr.org/archive/crypto2001/21390001.pdf
[50] http://dx.doi.org/10.1007/978-3-642-33125-1_11
[51] http://research.microsoft.com/en-us/projects/t2/

which has proved highly useful in security settings, be reasoned about in a similar way?

Acknowledgements. We would like to thank Jonathan Katz for his initial involvement in organizing the seminar and Matthew Hammer for his help in preparing the full report.

References

1 Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Łukasz Mazurek. Secure Multiparty Computations on Bitcoin. 2014 IEEE Symposium on Security and Privacy, SP 2014, Berkeley, CA, USA, May 18–21, 2014, pp. 443–458, IEEE CS. http://dx.doi.org/10.1109/SP.2014.35

2 Marcin Andrychowicz, Stefan Dziembowski, Daniel Malinowski, and Łukasz Mazurek. Secure Multiparty Computations on Bitcoin. Cryptology ePrint Archive, Report 2013/784, 2013. https://eprint.iacr.org/2013/784

3 Oded Goldrcich and Ronen Vainish. How to solve any protocol problem – an efficiency improvement (extended abstract). In Carl Pomerance, editor, *Advances in Cryptology – CRYPTO'87*, volume 293 of *Lecture Notes in Computer Science*, pages 73–86. Springer Berlin Heidelberg, 1988.

4 Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):690–728, July 1991.

4.73 Programming Languages for Big Data (PlanBig)

Organizers: James Cheney, Torsten Grust, and Dimitrios Vytiniotis

Seminar No. 14511

Date: December 15–19, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.12.48

Participants: Umut A. Acar, Yanif Ahmad, Alexander Alexandrov, Carsten Binnig, Giuseppe Castagna, James Cheney, Laurent Daynès, Nate Foster, Pierre Geneves, Alexey Gotsman, Todd J. Green, Torsten Grust, Fritz Henglein, Jan Hidders, Christoph Koch, Tim Kraska, Sam Lindley, Todd Mytkowicz, Thomas Neumann, Frank Neven, Ryan R. Newton, Kim Nguyen, Klaus Ostermann, Christopher Ré, Tiark Rompf, Andrew Stevenson, Julia Stoyanovich, Dan Suciu, Jens Teubner, Alexander Ulrich, Jan Van den Bussche, Stijn Vansummeren, Marcos Vaz Salles, Jan Vitek, Dimitrios Vytiniotis, Ryan Wisnesky

Large-scale data-intensive computing, commonly referred to as "Big Data", has been influenced by and can further benefit from programming languages ideas. The MapReduce programming model is an example of ideas from functional programming that has directly influenced the way distributed big data applications are written. As the volume of data has grown to require distributed processing potentially on heterogeneous hardware, there is need for effective programming models, compilation techniques or static analyses, and specialized language runtimes. The motivation for this seminar has been to bring together researchers working on foundational and applied research in programming languages but also data-intensive computing and databases, in order to identify research problems and opportunities for improving data-intensive computing.

To this extent, on the database side, the seminar included participants who work on databases, query languages and relational calculi, query compilation, execution engines, distributed processing systems and networks, and foundations of databases. On the programming languages side, the seminar included participants who work on language design, integrated query languages and meta-programming, compilation, as well as semantics. There was a mix of applied and foundational talks, and the participants included people from universities as well as industrial labs and incubation projects.

The work that has been presented can be grouped in the following broad categories:

- Programming models and domain-specific programming abstractions (Cheney, Alexandrov, Vitek, Ulrich). How can data processing and query languages be integrated in general purpose languages, in type-safe ways and in ways that enable traditional optimizations and compilation techniques from database research? How can functional programming ideas such as monads and comprehensions improve the programmability of big data systems? What are some language design issues for data-intensive computations for statistics?
- Incremental data-intensive computation (Acar, Koch, Green). Programming language support and query compilation techniques for efficient incremental computation for data set or query updates. Efficient view maintainance.
- Interactive and live programming (Green, Vaz Salles, Stevenson, Binnig, Suciu). What are some challenges and techniques for interactive applications. How to improve the live programming experience of data scientists? Ways to offer data management and analytics as cloud services.
- Query compilation (Neumann, Henglein, Rompf, Ulrich). Compilation of data processing languages to finite state automata and efficient execution. Programming languages techniques, such as staging, for enabling implementors to concisely write novel compilation schemes.
- Data programming languages and semantics (Wisnesky, Vansummeren). Functorial semantics for data programming languages, but also foundations for languages for information extraction.
- Foundations of (parallel) query processing (Suciu, Neven, Hidders). Communication complexity results, program equivalence problems in relational calculi.
- Big data in/for science (Teubner, Stoyanovich, Ré). Challenges that arise in particle physics due to the volume of generated data. Howe we can use data to speed up new material discovery and engineering? How to use big data systems for scientific extraction and integration from many different data sources?
- Other topics: architecture and runtimes (Ahmad), coordination (Foster), language runtimes (Vytiniotis), weak consistency (Gotsman).

The seminar schedule involved three days of scheduled talks,

followed by two days of free-form discussions, demos, and working groups. This report collects the abstracts of talks and demos, summaries of the group discussion sessions, and a list of outcomes resulting from the seminar.

4.74 Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis

Organizers: Jane Hillston, Jeremy Pitt, Martin Wirsing, and Franco Zambonelli

Seminar No. 14512

Date: December 14–19, 2014 | Dagstuhl Seminar
Full report – DOI: 10.4230/DagRep.4.12.68

Participants: Jacob Beal, Lenz Belzner, Luca Bortolussi, Giacomo Cabri, Rocco De Nicola, Giovanna Di Marzo Serugendo, Vashti Galpin, Marco Gribaudo, Salima Hassas, Jane Hillston, Annabelle Klarl, Roberta Lanciani, Peter Lewis, Michele Loreti, Stefano Mariani, Mieke Massink, Ugo Montanari, Laura Nenzi, Jeremy Pitt, Christophe Scholliers, Hella Seebach, Nikola Serbedzija, Mirco Tribastone, Petr Tuma, Danny Weyns, Martin Wirsing, Franco Zambonelli

Modern systems are often structured as complex, multi-layered networks of interconnected parts, where different layers interact and influence each other in intricate and sometimes unforeseen ways. It is infeasible for human operators to constantly monitor these interactions and to adjust the system to cope with unexpected circumstances; instead systems have to adapt autonomously to dynamically changing situations while still respecting their design constraints and requirements. Because of the distributed and decentralized nature of modern systems, this usually has to be achieved by collective adaptation of the nodes comprising the system. In open systems exhibiting collective adaptation, unforeseen events and properties can arise, e.g. as side effects of the interaction of the components or the environment. Modelling and engineering collective adaptive systems (CAS) has to take into account such "emergent" properties in addition to satisfying functional and quantitative requirements.

Finding ways to understand and design CAS, and to predict their behaviour, is a difficult but important endeavour. One goal of this seminar was to investigate techniques for modelling and analysing systems that adapt collectively to dynamically changing environment conditions and requirements. In many cases, these models and analysis techniques should not only capture qualitative properties of the system, such as absence of deadlocks, they should also be able to express quantitative properties such as quality of service.

Research on CAS builds on and integrates previous research efforts from several areas:

- Formal foundations and modelling techniques for concurrent systems deal with problems such as enabling and limiting concurrency, access to shared resources, avoidance of anomalies, communication between processes, and estimation of performance.
- Analysis of concurrent systems typically exploits such notions as bisimilarity of different processes or reasons on stochastic properties of systems consisting of many equivalent processes.
- The area of adaptive systems also investigates systems consisting of interacting entities, but is more concerned with the reaction of whole systems or individual actors in a system to a changing environment.

An important aim of this seminar was to combine research from concurrent systems with results from the adaptive systems community in order to develop formalisms for specifying CAS, to increase the scalability of qualitative and quantitative modelling and analysis techniques to large systems, and to apply them to systems that dynamically change their structure or adapt to novel situations.

The seminar was organised with a mixture of talks and working group sessions which facilitated more in-depth discussions and exploration of topics. In this report we include the abstracts of a selection of the presented talks, and three longer contributions compiled after the meeting which seek to reflect the activities of the working groups. The first group, considering modelling, specification and programming for CAS, start their presentation with brief descriptions of four diverse applications developed on the basis of CAS, ranging from national level power management to personal wearable devices. To complement this identification of application domains, the group also catalogued common and contrasting features that can be found in CAS. This consideration highlights the role of physical space in all the considered domains and the urgent need to develop modelling and analysis techniques which reflect this central role played by space. This was key amongst a number of challenges identified by the group in their conclusions. Spatio-temporal aspects were also identified as a key challenge by the second working group who considered verification of CAS. The report from this group outlines the role of verification within the design and management of CAS ranging

from seeking to guarantee global emergent behaviour from local specifications to using online verification to drive adaptation. Two specific challenges were explored in more detail, namely handling the inherent uncertainty in CAS, and specification and verification of spatial properties of systems composed of self-organising patterns. The third working group focused on the issues that arise from the recognition that some of the entities within a CAS may be humans and outside technological control, i.e. the design of socio-technical systems. A number of different scenarios are provided to illustrate the difference between socio-technical CAS and 'technical' CAS, and the human factors which must be taken into account. To remediate some of the problems identified, the group propose the idea of a general intervention framework, based around the 3I life-cycle – inspection-innovation-intervention. It was foreseen that intervention would be achieved by shaping mechanisms, and the report goes on to describe some possible shaping mechanisms which were considered. To conclude a number of research challenges are discussed.

5 Öffentlichkeitsarbeit
Public Relations and Outreach

5.1 Pressemitteilungen und Medienarbeit

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[52] 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. 660 Abonnenten verbreitet. Zunehmend nutzen aber auch Seminarteilnehmer den Dienst, um ihre Eindrücke vom Seminar mitzuteilen. Darüber hinaus werden über den Twitter-Nutzer @dblp_org Informationen über die Bibliographiedatenbank dblp verbreitet. Bei LinkedIn wird eine eigene Gruppe „Friends of Schloss Dagstuhl“ unterhalten (derzeit etwa 630 Mitgliedern), mit dem Ziel, die Vernetzung der Teilnehmer von Dagstuhl-Seminaren zu unterstützen. Weiterhin werden dort interessante Neuigkeiten rund um Schloss Dagstuhl bekannt gegeben.

5.1 *Press Releases and Media Work*

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[52].

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 660 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. At LinkedIn, a “Friends of Schloss Dagstuhl” group is maintained (with about 630 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.

5.2 Fortbildung

Schloss Dagstuhl engagiert sich im schulischen Bereich durch Organisation einer jährlichen Lehrerfortbildung, die sich an Informatiklehrer im Saarland und Rheinland-Pfalz richtet. Die Veranstaltung wird in Zusammenarbeit mit dem saarländischen Landesinstitut für Pädagogik und Medien (LPM) und dem Pädagogischen Landesinstitut RheinlandPfalz (PL) organisiert. Das Interesse an dieser Fortbildung stieg seit dem Beginn in 1991 stetig an und die 24. „Lehrerfortbildung in Informatik“', die vom 10. bis 12. Dezember 2014 statt fand, führte mehr Teilnehmer zusammen als jemals zuvor. Die intensive Fortbildung richtet sich zwar hauptsächlich an Lehrer aus dem Saarland und Rheinland-Pfalz, jedoch häufen sich Anfragen zur Teilnahme von Lehrern aus anderen Bundesländern. Mehr Informationen zur Veranstaltung 2014 gibt es auf der Webseite der Veranstaltung[53].

Um junge Journalisten und Volontäre zu ermutigen, über anspruchsvolle Informatikthemen zu berichten, hat

5.2 *Educational Training*

Schloss Dagstuhl holds an annual teacher training workshop specifically designed for teachers of 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). Interest in the workshop has risen steadily since the program began in 1991 and the 24th annual Dagstuhl Teacher Training Workshop, held at Schloss Dagstuhl on December 10–12, 2014, attracted more participants than ever before. While this intensive training program mainly targets teachers from the Saarland and the Rhineland Palatinate, Schloss Dagstuhl does receive requests for participation from teachers of other federal states. Details about the training in 2014 are available at the event webpage[53].

In order to encourage young journalists and trainees to report on complex informatics topics, Schloss Dagstuhl offered once again a workshop on science journalism. In

[52] http://www.dagstuhl.de/de/ueber-dagstuhl/presse/

Schloss Dagstuhl wieder einmal einen Workshop Wissenschaftsjournalismus angeboten. In 2014 fand er vom 25. bis 28. Mai statt. Als Dozenten für den Workshop konnten Tim Schröder (Wissenschaftsjournalist und Medientrainer, Oldenburg) und Gordon Bolduan (Pressesprecher des Excellenz-Cluster „Multimodal Computing and Interaction“ an der Universität des Saarlandes) gewonnen werden. Alle Teilnehmer als auch die Dozenten waren höchst zufrieden mit den Inhalten und Ergebnissen des Workshops. Weitere Informationen sind auf der Webseite des Workshops[54] abrufbar.

2014, the workshop took place on May 25–28. Trainers included Tim Schröder from Oldenburg (scientific writer and media trainer) and Gordon Bolduan (press relations officer at the Cluster of Excellence “Multimodal Computing and Interaction” at Saarland University). Participants as well as trainers and referees were very satisfied with the workshop. See the event webpage[54] for further details.

5.3 Wissenschaftsjahr 2014 – Die digitale Gesellschaft

5.3 *Year of Science 2014 – The Digital Society*

Zusätzlich zu den Fortbildungsmöglichkeiten beteiligte sich Schloss Dagstuhl an der interaktiven Ausstellung „Digital – der Fingerzeig nach vorn“, die vom 25. September bis zum 24. Oktober 2014 im Rahmen der durch den Bund geförderten Initiative „Wissenschaftsjahr 2014 – Die digitale Gesellschaft“ im Hauberisser Saal des Rathauses St. Johann in Saarbrücken stattfand. Das Ziel der Ausstellung war, Menschen aller Generationen anzuregen, sich an der Diskussion über die Entwicklung der „Digitalen Gesellschaft“ zu beteiligen. In Workshops und Gesprächsrunden konnten sich Interessierte über technologische Entwicklungen und gesellschaftliche Veränderungen, die damit verbunden sind, austauschen. Unter anderem konnten ältere Internet-Neulinge beim „Offenen Internet-Treff für Seniorinnen und Senioren“ ihre Angst vorm Computer verlieren und den Umgang mit dem Internet lernen.

Die zwölf Exponate der Ausstellung bezogen sich auf bekannte Phänomene und Alltagssituationen. Die Exponate zeigten, was heute schon möglich ist oder in Zukunft möglich sein wird. Mitarbeiter vor Ort erklärten den Interessierten, wie die Exponate funktionieren und welche Forschung dahinter steckt. Ganz nebenbei und spielerisch konnten die Besucherinnen und Besucher dabei einige Grundlagen der Informatik kennen lernen.

Beteiligte Forschungsinstitute und Kooperationen:

- Kompetenzzentrum für Informatik
- Max Planck Institute für Informatik und Softwaresysteme
- Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI)
- DFKI Spin-off white c
- Leibniz-Institut für Informatik Schloss Dagstuhl
- Landesmedienanstalt
- Staatskanzlei

In addition to the educational training opportunities, Schloss Dagstuhl also participated in the interactive exhibit “Digital – The Cue Forward” from September 25 to October 24, 2014, as part of the federally-sponsored “Year of Science 2014: Digital Society” initiative. The exhibition took place in the Hauberisser hall of the town hall of St. Johann in Saarbrücken. The aim of the exhibition was to encourage people of all generations for participating in the debate on the development of “Digital society”. In workshops and discussion groups, interested parties had the possibility to discuss technological developments and social changes that are associated with these developments. Among others, older Internet newbies were invited using the “Open Internet-Treff for senior citizens” to lose their fear of computers and learn how to use the Internet.

The twelve exhibits related to well-known phenomena and everyday situations. The exhibits showed what is possible today or will be possible in future. Site staff told the visitors how the exhibits operate and what kind of research is behind it. By the way and playful the visitors were able to get to know some of the basics of computer science.

Participating research institutions and cooperations:

- Kompetenzzentrum für Informatik
- Max Planck Institute für Informatik und Softwaresysteme
- Deutsches Forschungszentrum für Künstliche Intelligenz GmbH (DFKI)
- DFKI Spin-off white c
- Leibniz-Institut für Informatik Schloss Dagstuhl
- Landesmedienanstalt
- Staatskanzlei

[53] http://www.dagstuhl.de/14503
[54] http://www.dagstuhl.de/14222

Jahrestagung der Leibniz-Gemeinschaft

5.4

Annual Meeting of the Leibniz Association

Die Jahrestagung 2014 der Leibniz-Gemeinschaft fand wieder einmal in der letzten Novemberwoche in Berlin statt. Als Ort für die Eröffnung der Jahrestagung am Mittwoch, 26. November 2014, wurden die Räumlichkeiten der saarländischen Vertretung des Bundes in Berlin ausgewählt. Dies war der Anlass für die Staatskanzlei des Saarlandes, den beiden saarländischen Leibniz-Instituten Schloss Dagstuhl – Leibniz-Zentrum für Informatik und INM – Leibniz-Institut für Neue Materialien, die Möglichkeit zu geben, sich mit einem Stand zu präsentieren.

Neben einem Infoterminal, auf dem aktuelle Informationen über Dagstuhl und das Programm zu sehen waren, haben wir ausführliches Informationsmaterial über Schloss Dagstuhl, aber auch Ergebnisse aus Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops zum Mitnehmen zur Verfügung gestellt. Das Highlight des Abends war allerdings die in der Schlossküche hergestellte Marillenmarmelade, die von vielen Gästen, darunter die saarländische Ministerpräsidentin Annegret Kramp-Karrenbauer und der Leibniz-Präsident Prof. Matthias Kleiner, mehrfach lobend erwähnt wurde.

The annual meeting 2014 of the Leibniz Association was once again held in the last week of November in Berlin. For the opening of the annual meeting on Wednesday, November 26, 2014, the premises of the Saarland Representation of the Federal in Berlin was selected. This was the reason for the State Chancellery of the Saarland for giving the two Leibniz institutes of the Saarland, namely Schloss Dagstuhl – Leibniz Center for Computer Science and INM – Leibniz Institute for New Materials, the opportunity to represent themselves with an own booth.

In addition to an information terminal where we presented the latest information on Dagstuhl and its program, we offered extensive information to take away about Schloss Dagstuhl but also results from Dagstuhl Seminars and Dagstuhl Perspectives Workshops like Dagstuhl Manifestos. The highlight of the evening however was the apricot jam produced in the castle kitchen that was repeatedly praised by many guests including the Saarland Prime Minister Annegret Kramp-Karrenbauer and the Leibniz President Prof. Matthias Kleiner.

Fig. 5.1
Scientific Staff member Dr. Roswitha Bardohl at the 2014 annual meeting of the Leibniz Association. | Photo © Leibniz-Gemeinschaft/ Peter Himsel

6 Dagstuhl Publishing
Dagstuhl Publishing

Dagstuhl Publishing

Portfolio 6.1 *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.

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.

Dagstuhl Reports

Alle Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops werden in der Zeitschrift *Dagstuhl Reports*[55] dokumentiert, was eine Zitation der Seminare im wissenschaftlichen Kontext ermöglicht. Zudem erlaubt es auch denjenigen Wissenschaftlern, die nicht am Seminar teilgenommen haben, einen zeitnahen Einblick in das, was beim Seminar diskutiert und erarbeitet wurde.

Die Zeitschrift wurde 2011 ins Leben gerufen und enthält in monatlichen Ausgaben Berichte zu den Seminaren und Perspektiven-Workshops, die im jeweiligen Monat stattgefunden haben. Der Inhalt der Berichte wird nicht begutachtet. Das wissenschaftliche Direktorium (siehe Fig. 11.4) agiert als Herausgebergremium für die Reihe. Um umfassende Zusammenstellungen von begutachteten Artikeln auf Basis eines Dagstuhl-Seminars oder -Perspektiven-Workshops zu ermöglichen, wurde die Buchreihe *Dagstuhl Follow-Ups* (siehe unten) gegründet.

In 2014 wurden für alle 75 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*[55] which enables the citation of the seminars in a scientific context. Furthermore, it allows scientists who were not able to attend the seminar to inform themselves about the work and discussions of the seminar in a timely manner.

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

All of the 75 Dagstuhl Seminars and Dagstuhl Perspectives Workshops that took place in 2014 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*[56] die Manifestos der Dagstuhl-Perspektiven-Workshops – deren Erstellung zur Aufgabe des Workshops gehört – Open Access veröffentlicht. Das wissenschaftliche Direktorium (siehe Fig. 11.4) fungiert hier ebenfalls als Herausgebergremium. Die Ausgabe für 2014 enthält zwei Manifestos, siehe Fig. 6.1.

Dagstuhl Manifestos

Since 2011 we have published the manifestos – an expected result of Dagstuhl Perspectives Workshops – in the journal *Dagstuhl Manifestos*[56] in Open Access manner. The Scientific Directorate (see Fig. 11.4) acts as the editorial board of the journal. The 2014 volume includes two Dagstuhl Manifestos; see Fig. 6.1.

[55] http://drops.dagstuhl.de/dagrep
[56] http://drops.dagstuhl.de/dagman

Massive Open Online Courses: Current State and Perspectives
http://dx.doi.org/10.4230/DagMan.4.1.1
based on Dagstuhl Perspectives Workshop 14112 14112

Co-Design of Systems and Applications for Exascale
http://dx.doi.org/10.4230/DagMan.4.1.28
based on Dagstuhl Perspectives Workshop 12212

Fig. 6.1
Manifestos published in the 2014 volume of the journal *Dagstuhl Manifestos*.

Dagstuhl Publishing

Dagstuhl Follow-Ups

Die Buchreihe *Dagstuhl Follow-Ups*[57] ermöglicht die Veröffentlichung einer Sammlung begutachteter Beiträge, die auf einem Dagstuhl-Seminar oder Dagstuhl-Perspektiven-Workshop basiert. Für jedes Buch ist ein Antrag notwendig, der vom wissenschaftlichen Direktorium (welches als Herausgebergremium verantwortlich ist) begutachtet und freigegeben werden muss. In 2014 wurde kein Buch in der Reihe veröffentlicht.

OASIcs: OpenAccess Series in Informatics

Die *OASIcs*-Reihe[58] veröffentlicht begutachtete Tagungsbände von Workshops, Symposien und Konferenzen. Das Herausgebergremium (Fig. 6.2), 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 2014 wurden 8 Bände von thematisch breit gestreuten Workshops und Konferenzen veröffentlicht, siehe Fig. 6.3.

Dagstuhl Follow-Ups

The *Dagstuhl Follow-Ups*[57] 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 2014, no volume was published in the series.

OASIcs: OpenAccess Series in Informatics

The *OASIcs* series[58] aims to publish the peer-reviewed proceedings of workshops, symposia, and conferences. The editorial board, see Fig. 6.2, 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 2014, Dagstuhl published 8 *OASIcs* volumes covering the proceedings of topically widespread workshops and conferences; see Fig. 6.3.

[57] http://drops.dagstuhl.de/dfu
[58] http://drops.dagstuhl.de/oasics

Fig. 6.2
OASIcs Editorial Board.

Vol. 36 | MCPS'14 | 5th Workshop on Medical Cyber-Physical Systems
http://www.dagstuhl.de/dagpub/978-3-939897-66-8

Vol. 37 | SCOR'14 | 4th Student Conference on Operational Research
http://www.dagstuhl.de/dagpub/978-3-939897-67-5

Vol. 38 | SLATE'14 | 3rd Symposium on Languages, Applications and Technologies
http://www.dagstuhl.de/dagpub/978-3-939897-68-2

Vol. 39 | WCET'14 | 14th International Workshop on Worst-Case Execution Time Analysis
http://www.dagstuhl.de/dagpub/978-3-939897-69-9

Vol. 40 | WPTE'14 | First International Workshop on Rewriting Techniques for Program Transformations and Evaluation
http://www.dagstuhl.de/dagpub/978-3-939897-70-5

Vol. 41 | CMN'14 | 2014 Workshop on Computational Models of Narrative
http://www.dagstuhl.de/dagpub/978-3-939897-71-2

Vol. 42 | ATMOS'14 | 14th Workshop on Algorithmic Approaches for Transportation Modelling, Optimization, and Systems
http://www.dagstuhl.de/dagpub/978-3-939897-75-0

Vol. 43 | ICCSW'14 | 2014 Imperial College Computing Student Workshop
http://www.dagstuhl.de/dagpub/978-3-939897-76-7

Fig. 6.3
OASIcs volumes published in 2014.

Dagstuhl Publishing

LIPIcs: Leibniz International Proceedings in Informatics

Die *LIPIcs*-Reihe[59] veröffentlicht Tagungsbände von international renommierten Informatik-Konferenzen, die in ihrem jeweiligen Gebiet führend sind. Das internationale Herausgebergremium besteht aus einschlägig bekannten Wissenschaftlern und wird von Pascal Weil als Hauptherausgeber geleitet. Die Amtszeiten von Chris Hankin und Deepak Kapur wurden 2014 nach einem anonymen Wahlverfahren innerhalb des Herausgebergremiums bis Juni 2018 verlängert. Siehe auch Fig. 6.4.

In 2014 wurden Tagunsbände von fünf Konferenzen veröffentlicht: APPROX/RANDOM, FSTTCS, STACS, TYPES und TQC.

Die RTA-Konferenz wurde vom Herausgebergremium im Rahmen der bestehenden Kooperation erneut evaluiert und für weitere fünf Jahr (2015–2019) aufgenommen.

In 2014 gab so viele Bewerbungen bei LIPIcs wie nie zuvor, so dass für die kommenden Jahre ein deutlich gesteigertes Publikationsvolumen erreicht werden konnte. Die große Anzahl an Bewerbungen sind die erfreulichen Ergebnisse unserer langjährigen Bemühungen, einige der wichtigsten Konferenzen an LIPIcs zu binden. In Fig. 6.6 sind alle Konferenzen aufgelistet, deren Anträge bei LIPIcs positiv begutachtet wurden und mit denen daher eine mehrjährige Kooperation (typischweise 5 Jahre) eingegangen wurde. Verwaltungstechnisch werden nun Publikationsvereinbarungen mit den Steuerungsgremien der Konferenz abgeschlossen, um die gegenseitigen Verpflichtungen transparent zu machen und die Rollen zwischen Schloss Dagstuhl als Verlag und der Konferenz als Kunden festzulegen.

LIPIcs: Leibniz International Proceedings in Informatics

The *LIPIcs* series[59] 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; Pascal Weil acts as editor-in-chief. In 2014, the terms of Chris Hankin and Deepak Kapur were extended until June 2018, based on an anonymous voting within the editorial board. See also Fig. 6.4.

In 2014, the series published the proceedings of five major conferences: APPROX/RANDOM, FSTTCS, STACS, TYPES, and TQC.

RTA was re-evaluated by the LIPIcs editorial board and accepted for another five-year period (2015–2019).

Harvesting the fruits of our long-lasting efforts to attract major conferences to LIPIcs, the year 2014 has seen a substantial increase in applications for LIPIcs. Fig. 6.6 lists all conferences that have been accepted for a cooperation covering several years (typically 5 years). From the administrative side, publishing agreements are now signed between Dagstuhl and the steering committees of the conferences. The publishing agreement serves as a reference describing the commitments for Dagstuhl as a publisher and for the conference as a customer.

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

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

Fig. 6.4
LIPIcs Editorial Board.

Vol. 25 | STACS'14 | 31st International Symposium on Theoretical Aspects of Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-65-1
Vol. 26 | TYPES'13 | 19th International Conference on Types for Proofs and Programs http://www.dagstuhl.de/dagpub/978-3-939897-72-9
Vol. 27 | TQC'14 | 9th Conference on the Theory of Quantum Computation, Communication and Cryptography http://www.dagstuhl.de/dagpub/978-3-939897-73-6
Vol. 28 | APPROX/RANDOM'14 | Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques http://www.dagstuhl.de/dagpub/978-3-939897-74-3
Vol. 29 | FSTTCS'14 | 34th International Conference on Foundation of Software Technology and Theoretical Computer Science http://www.dagstuhl.de/dagpub/978-3-939897-77-4

Fig. 6.5
LIPIcs volumes published in 2014.

CALCO | Conference on Algebra and Coalgebra in Computer Science accepted for 2015–2020 (biennial conference, 3 editions)
CCC | Computational Complexity Conference accepted for 2015–2019
CONCUR | International Conference on Concurrency Theory accepted for 2015–2019
ESA | European Symposium on Algorithms accepted for 2016–2020
ICALP | International Colloquium on Automata, Languages, and Programming accepted for 2016–2020
ICDT | International Conference on Database Theory accepted for 2015–2019
IPEC | International Symposium on Parameterized and Exact Computation accepted for 2015–2019
SNAPL | Summit on Advances in Programming Languages accepted for 2015–2019 (biennial conference, 3 editions)
SoCG | Symposium on Computational Geometry accepted for 2015–2019
TLCA | International Conference on Typed Lambda Calculi and Applications accepted for 2015–2019

Fig. 6.6
Conferences that have have been accepted in 2014 for publication in LIPIcs for 2015–2020.

Dagstuhl Publishing

LITES: Leibniz Transactions on Embedded Systems

Die Open Access-Fachzeitschrift *LITES*[60] 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. 6.7), 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.

In 2014 wurde eine Absichtserklärung über die gemeinsame Herausgeberschaft mit der Fachgruppe *EMbedded Systems Special Interest Group (EMSIG)*[61] der Fachgesellschaft *European Design and Automation Association (EDAA)*[62] unterschrieben. Dabei ist 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 2014 wurden die ersten beiden Ausgaben von *LITES* mit insgesamt sechs Artikeln veröffentlicht.

LITES: Leibniz Transactions on Embedded Systems

The *LITES*[60] 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. 6.7), 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.

In 2014, a Memorandum of Understanding between Schloss Dagstuhl and the *EMbedded Systems Special Interest Group (EMSIG)*[61] of the *European Design and Automation Association (EDAA)*[62] regarding the joint publication of LITES was signed. 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 2014, the first two issues of *LITES* containing six articles in total has been published.

[60] http://drops.dagstuhl.de/lites
[61] http://www.emsig.net/
[62] http://www.edaa.com/

Dagstuhl Publishing

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. 6.7
LITES Editorial Board.

Infrastruktur 6.2 *Infrastructure*

Indizierung

Alle Reihen des Publikations-Portfolios werden bei *dblp* gelistet, siehe Fig. 6.8. 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. 6.8.

Die Bände der LIPIcs-Reihe waren zum Ende des Jahres 2014 bis inklusive Band 26 bei Scopus [63] indiziert. Zudem wurden im Berichtsjahr die technischen Schnittstellen verbessert, um die Datenakquisition von GoogleScholar zu unterstützen und somit die Publikationen sichtbarer und besser recherchierbar zu machen.

LeibnizOpen

Die Leibniz-Gemeinschaft hat mit *LeibnizOpen*[64] ein Online-Repositorium ins Leben gerufen, um Open Access-Veröffentlichungen von Leibniz-Instituten und deren Wissenschaftlern zu unterstützen und sichtbar zu machen. Schloss Dagstuhl liefert alle Artikel aus den Reihen *Dagstuhl Reports* und *Dagstuhl Manifestos* an das Repositorium und stärkt dadurch Forschungsergebnisse aus der Informatik innerhalb dieses multidisziplininären Repositoriums.

Indexing

All series of the publication portfolio are listed in *dblp*; see Fig. 6.8. 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. 6.8.

As of late 2014, LIPIcs volumes are now indexed in Scopus[63] up to volume 26. Furthermore, we have improved the technical interface of our publication server such that it can be harvested according to the guidelines of GoogleScholar. GoogleScholar now regularly retrieves metadata and full-texts from our server.

LeibnizOpen

The Leibniz Association has established the *LeibnizOpen*[64] 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.

[63] http://www.scopus.com
[64] http://www.leibnizopen.de/

DBLP	
Dagstuhl Reports	http://dblp.uni-trier.de/db/journals/dagstuhl-reports/
Dagstuhl Manifestos	http://dblp.uni-trier.de/db/journals/dagstuhl-manifestos/
Dagstuhl Follow-Ups	http://dblp.uni-trier.de/db/series/dfu/
OASIcs	http://dblp.uni-trier.de/db/series/oasics/
LIPIcs	http://dblp.uni-trier.de/db/series/lipics/

DOAJ	
OASIcs	http://doaj.org/toc/759bc28bcc174e25a1c571e9e29f9632
LIPIcs	http://doaj.org/toc/160b1ba80f8a46278ac8c92722c898c6
LITES	http://doaj.org/toc/ff079cbbc1744eeba13e3dd2ee16d937

Fig. 6.8
Indexing of Dagstuhl Publishing series in dblp and DOAJ.

Dagstuhl Publishing

AK Open Access der Leibniz-Gemeinschaft

Schloss Dagstuhl engagiert sich in der Arbeitsgruppe Open Access der Leibniz-Gemeinschaft. Im Rahmen dieses Engagements wurde ein Workshop „Erfolgreiches Journal-Management: Qualität und Reputation“ mit organisiert, welcher als Nachfolgeveranstaltung zu einem in 2013 erstmalig organisierten Workshop stattfand . Der Workshop fand am 30. und 31. Januar 2014 in der Geschäftsstelle der Leibniz-Gemeinschaft in Berlin statt. Der Workshop wurde von mehr als 50 Teilnehmer aus den Verlagsabteilungen von ungefähr 20 Leibniz-Instituten wahrgenommen. Für 2015 ist eine Nachfolgeveranstaltung mit dem Titel „Erfolgreiches Journal-Management: Sichtbarkeit und Strategie“ in Planung.

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“[65] 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)[66] werden alle Veröffentlichungen von Schloss Dagstuhl verwaltet. Es werden hierbei die allgemeinen Richtlinien für Online-Publikationen gemäß der Dublin Core-Initiative[67] berücksichtigt, wodurch alle nötigen Metadaten zu jeder Publikation gespeichert werden und die Langzeitverfügbarkeit sichergestellt wird. Die Online-Publikationen sind zitierfähig und stehen einer grossen Leserschaft zur Verfügung. Als technische Grundlage dient eine adaptierte Version des OPUS-Systems.[68]

Langzeitarchivierung

Alle Publikationen werden bei der Deutschen Nationalbibliothek (D-NB)[69] zur (digitalen) Langzeitarchivierung eingereicht.

Open Access Working Group of the Leibniz Association

A workshop entitled “Erfolgreiches Journal-Management: Qualität und Reputation” 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 30 and 31, 2014 and brought together more than 50 professionals in charge of publishing activities at about 20 Leibniz institutes. For 2015 a follow-up workshop entitled “Erfolgreiches Journal-Management: Sichtbarkeit und Strategie” is planned.

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.[65]

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)[66]. The general guidelines of the Dublin Core initiative[67] applicable to online publications are adhered to, meaning that all the requisite metadata of each publication is stored, thus ensuring availability in the long term. This enables the online publications to be cited by and accessible to a wide readership. The technical basis for this is an adapted version of the OPUS system.[68]

Long-term Archiving

All publications are submitted to the German National Library (D-NB)[69] for (digital) long-term archiving.

[65] http://www.allianzinitiative.de/de/handlungsfelder/open_access/
[66] http://www.dagstuhl.de/drops
[67] http://dublincore.org/
[68] http://elib.uni-stuttgart.de/opus/doku/about.php
[69] http://www.dnb.de/DE/Netzpublikationen/Langzeitarchivierung/langzeitarchivierung_node.html

Dagstuhl Publishing

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.[70] 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.[71]

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[70]. 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.[71]

[70] http://www.io-port.net (→ Digital Library → LIPIcs)
[71] http://vesta.informatik.rwth-aachen.de/Dagstuhl/

7 Bibliographiedatenbank dblp
dblp computer science bibliography

Offene Bibliographiedaten für die Informatik

7.1

Open Bibliographic Data in Computer Science

Moderne Informatik-Forschung benötigt den unmittelbaren und umfassenden Zugriff auf aktuelle Publikationen, um den Bedürfnissen in einer sich immer schneller entwickelnden und immer komplexer werdenden Forschungslandschaft gerecht zu werden. Doch nicht nur im 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.[72] Insgesamt ist dblp daher für die Informatik sowohl als Recherche-Tool, aber auch als Forschungsdatensatz unverzichtbar geworden.

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

[72] Google Scholar liefert zum Suchbegriff „dblp“ über 19 800 Treffer; im Einzelnen weisen SpringerLink ca. 1 800 Artikel, die ACM Digital Library ca. 1 500 Artikel, Elsevier ScienceDirect über 500 Artikel und IEEE Xplore über 160 Artikel nach.

Schloss Dagstuhl und dblp

7.2 Schloss Dagstuhl and dblp

Die Zusammenarbeit zwischen Schloss Dagstuhl und der ursprünglich an der Universität Trier entwickelten Bibliographiedatenbank dblp besteht bereits seit Ende 2010. Zunächst durch ein Projekt im Leibniz-Wettbewerb gefördert wird das Engagement seit Juni 2013 von Schloss Dagstuhl direkt mitfinanziert. Die Finanzierung wird zudem seit November 2010 durch eine großzügige Spende der Klaus-Tschira-Stiftung unterstützt. Bereits seit 2012 steht nun auch unter dblp.dagstuhl.de ein eigener dblp-Webservice unter der Domain von Schloss Dagstuhl bereit und ergänzt damit das dblp-Angebot der Universität Trier unter dblp.uni-trier.de. Das Kooperationsabkommen zwischen Schloss Dagstuhl und der Universität Trier wurde Ende 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. 7.1) leistet seit November 2011 unter dem Dach von Schloss Dagstuhl die wissenschaftliche Aufsicht und unterstützt das dblp-Team mit seiner Expertise.

The cooperation between Schloss Dagstuhl and the dblp computer science bibliography – originally developed at the University of Trier – has existed since late 2010. The commitment of Schloss Dagstuhl to dblp, initially funded by a project of the Leibniz Competition, has been funded directly by Schloss Dagstuhl since June 2013. Since November 2010, Schloss Dagstuhl's dblp team has also been supported by a generous donation from the Klaus Tschira Foundation. Schloss Dagstuhl's own dblp web service at dblp.dagstuhl.de was established in 2012 and complements the dblp service available at the University of Trier at dblp.uni-trier.de. In late 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 7.1), established in November 2011 at Schloss Dagstuhl, provides scientific supervision and supports dblp with its expertise.

[73] The search term "dblp" results in 19,800 hits at Google Scholar; in particular, SpringerLink lists 1,800 articles, the ACM Digital Libray lists 1,500 articles, Elsevier ScienceDirect lists more than 500 articles, and IEEE Xplore lits more than 160 articles.

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. 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. 7.1
dblp Advisory Board.

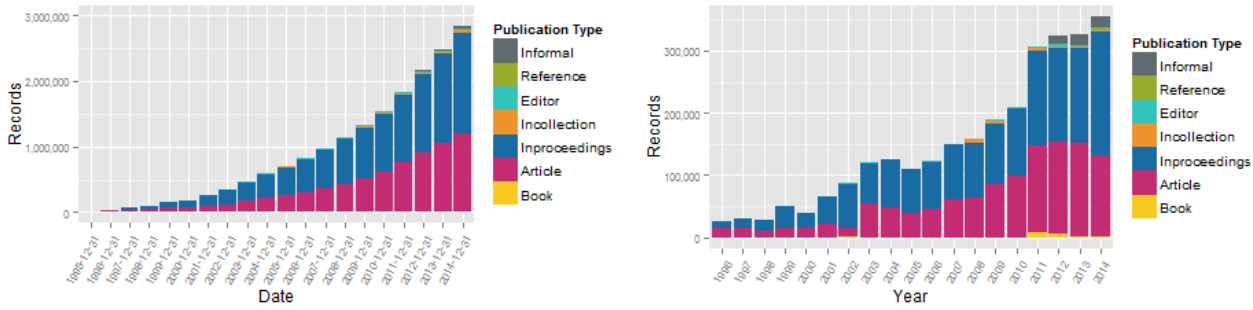

(a)Total number of records by year and type

(b)New records by year and type

Fig. 7.2
Development of the dblp data stock.

Statistiken der Datenakquise 7.3 *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 2014 indexierte dblp insgesamt 2 834 960 Publikationen. In dem Zeitraum von Anfang Januar 2014 bis Ende Dezember 2014 wurden dabei 354 437 neue Publikationseinträge in dblp aufgenommen. Diese Aufnahmequote stellt eine Steigerung um weitere 8,8 Prozentpunkte gegenüber der erst im Vorjahr erzielten Rekordaufnahmequote dar. Die neu aufgenommenen Einträge verteilen sich zu 55,6% auf Konferenzbeiträge, zu 37,1% auf Journalartikel, sowie zu 7,3% auf andere Publikationstypen.

Die Anzahl der in dblp erfassten Autoren überstieg am 18. Dezember 2014 die 1,5-Millionen-Marke. Ein Überblick über die Entwicklung der Datenakquise kann Fig. 7.2a und Fig. 7.2b 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, 2014, and December 31, 2014, the dblp database grew by 354,437 publication records to reach a total of 2,834,960 records. This number of new records is an 8.8 percent increase when compared to the already high rate of new inclusions reached in 2013. Of these new records, 55.6% have been conference papers, 37.1% have been journal articles, and 7.3% have been other publications.

The total number of authors indexed in dblp passed the 1,500,000 authors milestone on December 18, 2014. The development of the dblp dataset is summarized in Figure 7.2a and Figure 7.2b.

Nutzungsstatistiken 7.4 *Usage Statistics*

2014 wurden vom dblp-Team drei offizielle dblp-Server geführt. Die Daten dieser Server werden täglich aktualisiert und miteinander synchronisiert:

- Server Trier 1: http://informatik.uni-trier.de/~ley/db
- Server Trier 2: http://dblp.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

In 2014, three official dblp web servers were updated and synchronized on a daily basis:

- server Trier 1: http://informatik.uni-trier.de/~ley/db
- server Trier 2: http://dblp.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

	Trier 1	Trier 2	Dagstuhl
user sessions (visits) per day	21,057	4,703	326
page views per day	174,247	47,531	14,964
page views per user session	8.2	10.1	45.8
distinct users (IPs) per month	327,299	76,566	4,399
data served per month	825.2 GB	345.8 GB	27.2 GB
as above, including bots	5187.5 GB	672.2 GB	56.4 GB

Fig. 7.3
Average usage of the three dblp servers in 2014.

die mit Abstand bekannteste Adresse besitzt. Allerdings verzeichnen auch die kleineren dblp-Server einen stetig steigenden Besucherzuwachs. Zum Beispiel hat sich die Anzahl unterschiedlicher Nutzer (gezählt gemäß unterschiedlicher IPs) des Servers Dagstuhl im Laufe des Jahres 2014 von 3 235 auf 6 915 Nutzer annähernd verdoppelt.

Fig. 7.3 fasst die durchschnittliche Nutzung aller drei dblp-Server zusammen. Diese Statistiken ignorieren die Zugriffe, die durch bekannte Bot- und Crawler-Software verursacht wird. In der Tat ist der durch Bots und Crawler verursachte Datenverkehr inzwischen für etwa 80% des gesamten Datenvolumens des dblp-Dienstes verantwortlich.

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

engine. However, the smaller servers are catching up, as they have been hosting a beta build of the new dblp web system in the past year. For example, the number of distinct users (identified by IP) of the server Dagstuhl essentially doubled from 3,235 users in January 2014 to 6,915 users in December 2014.

Figure 7.3 shows the average usage of all three servers in 2014. These figures ignore the traffic caused by known bots and crawlers. In fact, the traffic caused by bots and crawlers accounts for over 80% of the data served by the dblp service in 2014.

Please note that these figures do not include the traffic of the (currently) still independently operated CompleteSearch front-end at dblp.org, which is maintained by the research group of Prof. Hannah Bast at the University of Freiburg. The dblp.org domain will be integrated into the dblp mirror structure in early 2015.

Gemeinsames Projekt von dblp, Zentralblatt MATH und HITS

7.5

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 dies Problem in seiner allgemeinen Form als das Problem der „Entity-Resolution“ oder der „Autorennamen-Disambiguierung“, und es stellt ein wichtiges Forschungsthema im Bereich der linguistischen Datenverarbeitung dar. 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 teils disjunkten Einträgen, stellt dabei eine interessante Möglichkeit dar, Fehler in den Datensätzen aufzudecken und von einander zu lernen. Die Natural-Language-Processing (NLP) Forschungsgruppe des HITS um Prof. Michael Strube bringt dabei ihre Erfahrung mit graph- und netz-

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 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

werkbasierten 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. Der Antrag erhielt durchweg positive Gutachten und die beantragte dreijährige Förderung wurde in vollem Umfang bewilligt. Der Leibniz-Senatsausschuss Wettbewerb (SAW) wertete das Projekt als *„einen Beitrag, die Qualität hinterlegter Metainformationen in digitalen Sammlungen wissenschaftlicher Veröffentlichungen in erheblichem Maße zu steigern“* und erkennt in der geplanten Methodik *„einen sinnvollen Lösungsansatz für die gegebene Problemstellung.“*

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 proposal received generally positive reviews and was accepted for 3 years of funding. The Leibniz Senate Competition Committee (SAW) has favorably described the project as *“a contribution to significantly improve the quality of metadata in digital collections of scientific publications”* and considers the proposed methodology *“a meaningful solution to the given problem.”*

Datenabruf für das CHE Hochschulranking

7.6

Data Extraction for the CHE German University Ranking

Im Juli 2014 erreichte Schloss Dagstuhl eine Anfrage von der Gütersloher Denkfabrik „Centrum für Hochschulentwicklung“ (CHE), einem von der Bertelsmann-Stiftung und der Hochschulrektorenkonferenz (HRK) gegründeten Institut. Das CHE veröffentlicht unter anderem in Kooperation mit Die Zeit regelmäßige Hochschulrankings zu verschiedenen Fachdisziplinen, darunter auch Informatik. Zuvor hatte sich der Informatik-Fachbeirat des CHE, bestehend aus Vertretern des Fakultätentag Informatik (FTI) und des Fachbereichstags Informatik (FBTI), für die Verwendung bibliographischer Analysen basierend auf Daten aus dblp im kommenden Informatik-Ranking ausgesprochen.

Das dblp-Team wurde gebeten, alle in dblp gespeicherten bibliographischen Daten der Professoren und Post-Doktoranden von 65 deutschen Informatik-Fakultäten zu extrahieren und an das CHE zu übersenden. Dazu wurden von den Fakultäten selbst entsprechende Namenslisten zur Verfügung gestellt. Im Rahmen eines Werkvertrages wurde der Datenabruf von einem Mitglied des dblp-Teams von Schloss Dagstuhl durchgeführt und die rohen Publikationsdaten im Frühjahr 2015 an das CHE übermittelt.

Das CHE beabsichtigt, die extrahierten Daten in Zusammenarbeit mit dem Team Bibliometrie der Zentralbibliothek am Forschungszentrum Jülich zu untersuchen. Ein zum Teil auf diesen Analysen basierendes Ranking der deutschen Informatik-Fakultäten soll im Laufe des Jahres 2015 in Die Zeit veröffentlicht werden. Schloss Dagstuhl und das dblp-Team sind dabei weder an den bibliometrischen Analysen noch an der Erstellung des Rankings beteiligt.

In July 2014, Schloss Dagstuhl was approached by the Gütersloh-based think tank “Center for Higher Education” (CHE), an institute founded by the Bertelsmann Foundation and the German Rectors’ Conference (Hochschulrektorenkonferenz, HRK). CHE publishes a ranking of a number of German computer science faculties in cooperation with the German newspaper *Die Zeit*. This ranking is supported by an advisory committee consisting of representatives of the Fakultätentag Informatik (FTI) and the Fachbereichstags Informatik (FBTI).

The dblp team of Schloss Dagstuhl was asked to extract from dblp the bibliographic records of all professors and post-doctoral researchers for 65 German computer science faculties. Complete name lists of those researchers were provided by the faculties themselves; the dblp data extraction was executed by a member of the dblp team. The data was submitted to CHE in early 2015 as the subject of a service contract between CHE and Schloss Dagstuhl.

CHE plans to evaluate the extracted bibliographic data in corporation with the Bibliometrics Team of Forschungszentrum Jülich. A ranking is expected to be published in *Die Zeit* later in 2015. Schloss Dagstuhl and the dblp team are not involved in the interpretation of the data nor the creation of the ranking.

8 Einrichtung und Service
Facilities and Services

Freizeit und Ambiente 8.1 *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 zur Verfügung stehen. 2014 wurde der Bestand um zwei Konzertgitarren erweitert. Unser Zentrum verfügt außerdem über eine Sauna, einen Billardtisch, Tischfußball, Mountainbikes, eine Dartscheibe, einen Freizeitraum mit Fitnessgeräten und Tischtennis sowie einen Außenbereich mit Volleyballnetz und Bouleplatz.

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. The collection was expanded in 2014 with the acquisition of a second Höfner concert guitar. 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, and outdoor sports grounds featuring a net and boules area.

Kinderbetreuung 8.2 *Childcare*

Schloss Dagstuhl bietet allen Teilnehmern von Dagstuhl-Seminaren und Dagstuhl-Perspektiven-Workshops, die mit Kindern anreisen, ein qualifiziertes Betreuungsprogramm für Kinder an. Dieser Service kann im voraus gebucht werden. Alternativ ist es Eltern auch möglich eine Begleitperson zur Betreuung des Kindes oder der Kinder mitzubringen. Die Kosten für Verpflegung und Unterkunft für die Begleitperson als auch der Kinder übernimmt Schloss Dagstuhl.

Im Jahre 2014 wurden 22 Kinder auf Schloss Dagstuhl betreut. Davon wurden 9 Kinder durch einen Tagesmutter und 13 weitere durch Verwandte betreut. Positive Rückmeldungen von den Eltern sowie Dankeschön-Briefe und Bastelarbeiten, die die kleinen Gäste 2014 für das Kindergästebuch von Schloss Dagstuhl angefertig haben, sind in Kapitel 3 und über den ganzen Bericht verteilt zu finden.

Schloss Dagstuhl gladly offers to organize child care with a certified nanny for participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops who need to visit our center with young children. The service, which supports families and particularly women computer scientists, can be booked in advance of the seminar.

Parents also have the option to bring along their own "nanny," usually a spouse or relative, whose room and board costs are absorbed by the center just as they are for children. In 2014, Dagstuhl hosted 22 children, 9 of whom were cared for by a nanny on site and 13 by relatives. Positive feedback from parents, and thank-you notes and artwork created by center's youngest guests in 2014 for the Dagstuhl children's guest book, can be found in Chapter 3 and throughout this report.

Dagstuhls Küche 8.3 *Dagstuhl's Kitchen*

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 kennen zu lernen. Große Tische im Speiseraum fördern die gemeinschaftliche Interaktion bei den Mahlzeiten. Während der Pausen werden vor den Vortragsräumen heiße Getränke und Gebäck auf einem Kaffeewagen angeboten. Darüber hinaus gibt es in der „alten" Cafeteria und dem Weinkeller den ganzen Tag über warme Getränke und kleine Snacks. Abends gibt es in diesen beiden beliebten Räumen Brot und eine Käseauswahl.

Dagstuhls Philosophie des Kochens ist einfach: saisonal, gesund und schmackhaft. Unsere Gerichte werden

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. In addition to strategically-placed coffee and snacks during scheduled breaks, guests can enjoy a warm drink and a snack at any time in the café or the wine cellar – two traditional after-hours hangouts. Bread and cheese is served in both rooms every evening.

The philosophy behind Dagstuhl's cooking is simple: seasonal, healthy and tasty meal. Everything is freshly prepared each day by the kitchen's 10-person staff and apprentices in training. The focus is on lighter fare

jeden Tag frisch von unseren 10 Mitarbeitern der Küche und unserer Auszubildenden zubereitet. Der Schwerpunkt liegt dabei auf leichtem Essen während des Tages, um unsere Gäste nicht zu ermüden, und auf warmen Gerichten am Abend. Dieses steht ein wenig im Widerspruch zur deutschen Tradition, kommt aber der Mehrheit der internationalen Gäste des Zentrums durchaus entgegen. Hausgemachter Kuchen, Kaffee und Tee unterbrechen angenehm die tägliche Routine und anstrengenden Diskussionen.

Sowohl die Zutaten als auch die Gerichte wechseln entsprechend der Saison. An warmen Sommerabenden wird häufig auf der Terrasse vor dem Speisesaal gegrillt und unter anderem saarländische *Schwenker* (eine lokale Variante des Grillsteaks, die unter dauerndem Schwenken des Grillrostes zubereitet wird) den Gästen angeboten. 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. Auf Nachfrage bereiten unsere Mitarbeiter auch abgestimmte Gerichte für Vegetarier, Veganer und Gäste mit Lebensmittelunverträglichkeiten zu. 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 seit dem Ende 2014 verpflichtet sind. Des weiteren achten wir auf deklarationsfreie Zusatz- und Konservierungsstoffe.

Um unseren Gästen trotz eines beschränkten Budgets eine ausgewogene Qualität anbieten zu können, bietet unsere Küche ein Frühstücksbüffet, ein Mittagsbuffet sowie ein Menü am Abend an. Unser Restaurant mit den großen Fenstern zum Garten des Hauptgebäudes bietet ca. 80 Personen Platz. Hier herrscht eine entspannte und fast familiäre Atmosphäre, was nicht zuletzt auf unsere freundlichen und engagierten Mitarbeiter zurückzuführen ist.

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 internationality of the center's guests. Homemade cake, coffee and tea each afternoon punctuate the daily routine.

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. Special dishes for those with medical food conditions and vegetarians or vegans are prepared upon request. The kitchen works in accordance with the HACCP-Konzept (Hazard Analysis and Critical Points Concept) and adheres to the mandatory labeling of allergens to which all food processing establishments are required as of December 2014. Food additives and conservatives for which labeling is non-mandatory are also carefully monitored.

To accomplish all of this within a reasonable budget, the center offers a buffet-style breakfast and lunch and a set evening meal served by the kitchen's friendly and engaged staff. Guests find at Dagstuhl a relaxed, family-style atmosphere in the large, 80-person dining hall, which opens onto the Schloss garden with its relaxing patio.

8.4 Tagungsräume / *Conference Facilities*

Schloss Dagstuhl bietet drei Hörsäle für 25 bis 60 Personen. Alle Hörsäle sind mit einem Beamer, einen MS-Windows-Arbeitsplatz und einer Audioanlage einschließlich Mikrophone ausgestattet. Durch diese Technik werden Vorträge, Präsentationen und Live-Vorführungen auch verteilter Systeme optimal unterstützt. Mittels einem Presenter können Vortragende ihre vorbereiteten Materialien präsentieren, ohne zum Laptop oder Arbeitsplatz zurückzukehren.

Neben den Hörsälen bietet Dagstuhl sechs Seminarräume. Davon sind zwei mit modernen HDMI-fähigen Beamern ausgestattet, während in einem Hörsaal ein großes Plasmadisplay montiert ist. Fünf Beamer auf Rollwagen stehen zusätzlich zur flexiblen Benutzung in 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 sogar eine ganze Wand als Whiteboard (über $12\mathrm{m}^2$) benutzt werden, da diese mit einer speziellen Farbe gestrichen wurde.

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 workplace, 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 HDMI 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

Daneben gibt es eine Anzahl weiterer Orte, an denen Gäste sich zur Diskussion in entspannter Atmosphäre treffen können. 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.

8.5 Computer und Vernetzung

Schloss Dagstuhl bietet seinen Gästen eine adäquate Anbindung an das Internet. Seit 2013 erfolgt die Anbindung an das Internet über das DFN mit zwei 100 Mbit/s Leitungen. Bisher wurde einer dieser Leitung als sogenannten Backup- bzw. Failover-Leitung betrieben. Seit 2014 werden beide Leitungen nun gleichzeitig betrieben und der gesamte Netzwerkverkehr auf die zwei Leitungen verteilt und zwar so, dass den Gästen eine dieser Leitungen vollständig und exklusiv zur Verfügung steht. Fast 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. Die Seminar-Organisatoren haben im größten Seminarraum „Saarbrücken“ die Möglichkeit, den WLAN Empfang abzuschalten.

Neben dem Zugang über mitgebrachte Laptops, Tablet Computer oder Smartphones stehen den Gästen einige fest installierte Arbeitsplätze zur Verfügung. Schloss Dagstuhl hat einen Rechnerraum mit sechs festen Arbeitsplätze. Davon sind zwei Arbeitsplätze mit Apple Macs ausgerüstet, drei sind dedizierte MS-Windows-Arbeitsplätze und ein weiterer Arbeitsplatz mit Ubuntu Linux. Zusätzlich steht Benutzern eines Laptops ein externer Monitor samt Tastatur und Maus zur Verfügung. In diesem Raum stehen auch direkte Ethernet-Anschlüsse zur Verfügung, um das WLAN wegen Bandbreite oder Kapazitätsgründen zu umgehen. Weiterhin bietet Schloss Dagstuhl seinen Gästen zwei iPads sowie auf Nachfrage einen MacBook Pro und einen Laptop mit MS Windows.

Im Zentrum steht den Gästen ein Multifunktions-Farbdrucker mit Scanner und Kopierer und ein weitere Farbdrucker zur Verfügung. Der Zugriff erfolgt vorzugsweise über eine Weboberfläche, die das direkte Drucken zahlreicher Dokumentenformate erlaubt. Alternativ können die Drucker mittels entsprechender Treiber auch direkt aus dem lokalen Netzwerk angesteuert werden. Zu der IT-Ausstattung gehören weiterhin fünf Recherche-Arbeitsplätze in der Bibliothek sowie drei fest installierte Rechner in den Seminarräumen.

Seit 2014 werden die täglichen Datensicherungen auf jeweils zwei unterschiedlichen Tape-Libraries gespeichert. Die zwei Tape-Libraries befinden sich in unterschiedlichen Gebäuden, so dass im Katastrophenfall (Brand, Überschwemmung, ...) eine komplette Datensicherung zur Verfügung steht. Eine Library steht im Serverraum (Neubau) und die andere befindet sich im Technikraum des neu erbauten Gästehauses. Verbunden sind beide Gebäude über eine Glasfaser-Leitung. Trotz der großen Entfernung

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 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.

8.5 *Computers and Networks*

Schloss Dagstuhl offers its guests an adequate connection to the Internet. Since 2013 the center is connected to the Internet by the DFN (German Research Network) using two redundant 100 Mbit/s lines. Throughout the grounds guests have Internet access by Wi-Fi (IEEE 802.11 b,g,n). Access is either via eduroam or a Dagstuhl-hosted private account. In our largest seminar room, “Saarbrücken,” organizers may choose to disable the main Wi-Fi connection during meeting times.

Most of our guests prefer to access the Internet via their laptops, tablet computers and smartphones, but they are also free to use the workstations in our computer room. Schloss Dagstuhl offers one computer room including eight workstations. Among them there are two Apple Macs, two dedicated MS Windows workstations and four workstations providing either Linux or MS Windows by a dual boot method. There is also a large display together with an external keyboard and mouse for users with their own laptop. Several Ethernet cables with Internet connection are also provided to bypass the rate-restricted Wi-Fi connection. Two iPads, and upon request a MacBook Pro and a laptop with Windows are also available for use throughout the grounds.

Schloss Dagstuhl provides a multifunction color printer with scanner and copier, a color printer, and a black and white printer. The preferable access method is to use a dedicated web front end which allows to upload and print the most used document formats without converting them. Alternatively, guests can use the appropriate printer drivers on their computers to directly access the printers via the network. The center’s IT equipment also includes five workstations in the library for literature research, as well as three fixed computers in the lecture halls.

As of 2014, daily data backups are now stored on two different tape libraries. The two tape libraries 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 available in the server room of the New Building, while the other is housed in the technical utilities room of the new Guest House. Both buildings are connected via a fiber-optic line. Despite the great distance (about 200m), it was possible to use a fiber channel connection for this, which reaches a speed of 4 Gbit/s.

(ca. 200m) ist es möglich für diese Verbindung eine Fiber-Channel-Verbindung zu nutzen, die eine Geschwindigkeit von 4 Gbit/s erreicht.

Dagstuhl's Web-basierte Dienste 8.6 *Dagstuhl's Web-based Services*

Schloss Dagstuhl bietet allen Organisatoren und Gästen eine wachsende Anzahl Web-basierter 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, die für alle anderen zugreifbar sind. 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.

Schloss Dagstuhls Internetauftritt[74] bietet nicht nur seinen Gästen sondern allen Nutzern weltweit Informationen über die folgenden Themen:

- Verbreitung allgemeiner Informationen über das Zentrum, wie Konzept, Programm, Antragsmodalitäten, Stiftung
- 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 – fast 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.

Im November 2014 wurde unsere Plattform zum Austausch von Seminarmaterialien so aktualisiert, dass ihr Aussehen dem unserer Webseite angepasst wurde. Die Funktionalität wurde dabei nicht geändert. Die neue Version steht allen Seminaren ab dem Jahr 2015 zur Verfügung.

Schloss Dagstuhl offers an increasing number of web-based services to seminar organizers and participants. During the preparation phase, the seminar organizers can check 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.

In keeping with the center's philosophy, its Internet[74] 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, the Foundation
- 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
- Provision of 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.

In November of 2014, the platform offered by Schloss Dagstuhl for the exchange of seminar-related materials was updated in order to adapt its appearance to that of the Schloss Dagstuhl website. The new version of the platform, for which the functionality remained unchanged, was made available to all Dagstuhl Seminars starting in 2015.

[74] http://www.dagstuhl.de/

9 Bibliothek
Research Library

Bestand und Angebot 9.1 *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 2014 umfasste der Bibliotheksbestand 60 723 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 im 1. Obergeschoss 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[75] sind u.a. der Online-Bibliothekskatalog, die Zeitschriftenbestandsliste

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, 2014, the library's collection totaled 60,723 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 which are displayed on the library's first floor.

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.

mit Zugang zu den abonnierten online verfügbaren Zeitschriften sowie weitere Informationsangebote der Bibliothek zu erreichen.

The online catalogue and a comprehensive journal list with access to the subscribed journals as well as other information offerings can be accessed via the library's webpage.[75]

Spenden an die Bibliothek 9.2

Die Bibliothek von Schloss Dagstuhl profitiert durch zahlreiche Spenden. So erhielt die Informatik-Fachbibliothek im Jahr 2014 Buchspenden von den Verlagen, die in Fig. 9.1 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 949 Bände als Spenden von Verlagen und Seminarteilnehmern.

Library Donations

The Dagstuhl Informatics Research Library receives numerous book donations from publishers and seminar participants. During 2014 the Informatics Research Library received book donations from the publishers listed in Fig. 9.1. The center is also grateful for donations of author's copies, particularly those of major works that are out of print. The center received a total of 949 volumes during the year 2014 as donations from publishing houses and seminar participants.

[75] 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. 9.1
Donations from publishers to the Dagstuhl library.

10 Kunst
Art

Dagstuhl als Galerie 10.1 *Dagstuhl as Art Gallery*

Im sogenannten Kreuzgang des Neubaus werden regelmäßig Kunstausstellungen organisiert. Das großzügige Raumangebot der Wände des 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 fünf Kunstausstellungen für jeweils zwei bis drei Monate. Die vier Ausstellungen, die im Jahr 2014 stattfanden, sind nachfolgend beschrieben.

Fabian Treiber, »Neun Minuten vor Vegas«

Die Ausstellung mit den Werken des Stuttgarter Künstlers Fabian Treiber fand von Januar bis März 2014 auf Schloss Dagstuhl statt. 1986 in Ludwigsburg geboren, absolvierte Fabian Treiber zunächst eine Ausbildung zum Informatiker, bevor er 2007 ein Studium an der Akademie der bildenden Künste Stuttgart aufnahm und 2014 abschloss. Baumeisters Suche nach dem „Unbekannten," bei welcher das angestrebte Ziel ausschließlich als Methode in der Formfindung einer autonomen Bildwelt seine Berechtigung hat, ist eine Voraussetzung für seine Arbeit.

In der Berichterstattung über die Ausstellung wurde Treibers Technik als Dialog zwischen Maler und Gemälde beschrieben, in dem sich der Künstler bewusst einer gewissen Zeitnot aussetze, indem er die noch feuchte Ölfarbe auf der Leinwand lackiere. „ So entwickelt sich der Dialog zwischen Maler und Bild. Er muss damit auf Prozesse reagieren, die ihn überraschen. Es ist seine persönliche Handschrift, die so immer wieder anders entsteht."[76]

Aloys Ohlmann (1938–2013), »Fensterbilder – Hommage à Oskar Schlemmer«

Bei der vom März bis Mai in Schloss Dagstuhl stattfindenden Gedächtnisausstellung für Aloys Ohlmann (†16.09.2013) wurde erstmals seine Serie von 36 Fensterbildern gezeigt, die 1991 entstand. Aloys Ohlmann schuf diese Werkreihe als Hommage an den Bauhauskünstler Oskar Schlemmer, der 1942 – im Jahr vor seinem Tod – einen Zyklus von 18 Bildern schuf, die Fensteraus- und -einblicke aus unterschiedlicher Sicht darstellen. Die erste Ausstellung von Arbeiten Aloys Ohlmanns in Dagstuhl fand 2001 statt.

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 five art exhibits per year, with each exhibit generally running for two to three months. The four exhibits hosted by Schloss Dagstuhl in 2014 are described below.

Fabian Treiber, »Neun Minuten vor Vegas«

Featuring the works of Stuttgart-based artist Fabian Trieber, this exhibit ran at Schloss Dagstuhl from January to March of 2014. Treiber, who was born in 1986 in Ludwigsburg, Germany, initially trained as a computer scientist before moving on to study at the Academy of Fine Arts in Stuttgart, where he graduated in 2014. Taking as a prerequisite Baumeister's search for the "unknown," in which the objective pursued is only meaningful as a method in the form-finding of an autonomous visual world, Treiber's work moves between abstraction and representation.

Press coverage of the exhibit described Treiber's method as a dialogue between the painter and the painting in which the artist imposes an artificial time constraint on his process, varnishing the oil paints while they are still wet. During the creative process, he reacts spontaneously to surprises that come up, making it impossible for him to predict how the end result will appear.[76]

Aloys Ohlmann (1938–2013), »Window Pictures – Hommage to Oskar Schlemmer«

This memorial exhibition in honor of Aloys Ohlmann († September 16, 2013), which ran at Schloss Dagstuhl from March to May 2014, displayed for the first time a series of 36 window paintings created by the artist in 1991. Ohlmann painted these works in homage to Bauhaus artist Oskar Schlemmer, who in 1942 – one year before his death – completed a cycle of 18 paintings representing different window views as seen from different vantage points. The first exhibit of Ohlmann's work at Schloss Dagstuhl took place in 2001.

[76] *Saarbrücker Zeitung* 21.01.2014, „Er lässt sich von seinen Bildern überraschen".

Aloys Ohlmann, 1938 im saarländischen Baltersweiler geboren, studierte an der Saarbrücker Werkkunstschule (heute Hochschule der Bildenden Künste Saar), an der er von 1967 bis 1970 auch lehrte, und an der Staatlichen Akademie der bildenden Künste Stuttgart. Für sein Werk wurden ihm mehrere Preise verliehen, darunter 1979 der „Jerg Ratgeb"-Publikumspreis der Stadt Stuttgart und 2003 der Mia Münster-Preis der Stadt St. Wendel. Außerdem war er 1968 Mitbegründer der Künstlervereinigung „Gruppe 7" und der Galerie im Zwinger in St. Wendel.

In der Berichterstattung über die Ausstellung wurde Aloys Holmann als „Streiter für eine bessere Welt, ein politisch engagierter Künstler" beschrieben. Prof. Karl Otto Jung, ein enger Freund des Künstlers, stellte den Gästen auf der Vernissage Ohlmanns Werk vor. „Sie werden sehen und erleben, wie jeder Blick durch ein Fenster Ihnen einen neue Sicht auf einen Aspekt von Welt öffnet, mal heiter, mal traurig, mal beängstigend, aber immer aktuell Bezug nehmend auf die Welt." Prof. Jung forderte die Gäste auf, Ohlmanns Worte „Ich denke mehr, als in Wirklichkeit ist" zu bedenken.[77]

Anne Haring und Michael Mahren, »Skulpturen und Zeichnungen«

Schloss Dagstuhl beherbergte diese Ausstellung mit Werken zweier Künstler, Skulpturen von Anne Haring und Zeichnungen von Michael Mahren, von Mai bis Juli 2014.

Anne Haring, 1961 in Hamburg geboren, studierte von 1980 bis 1986 an der Staatlichen Kunstakademie Düsseldorf bei Prof. Karl Bobek. Während ihres Studiums verbrachte sie als Stipendiatin auch einige Zeit an der Cité Internationale des Arts in Paris. Nach ihrer Lehrtätigkeit in Zeichnen und Fotografie von 1999 bis 2004 an der TU Dortmund arbeitet und lebt Anne Haring nun als freie Künstlerin in Saarbrücken. 1997 erhielt sie den Jakob-Felsing-Preis der Volksbank Saarbrücken.

Michael Mahren, Jahrgang 1948, schloss 1976 sein Studium an der Fachhochschule Saarbrückenbei Oskar Holweck und Robert Sessler als Grafikdesigner ab. 1980 erhielt er gemeinsam mit Axel Reiß den ersten Preis im Wettbewerb zur künstlerischen Gestaltung des Vorplatzes der Alten Feuerwache in Saarbrücken. Seit 2005 ist Michael Mahren Vorsitzender des BBK Saar.

Gemeinsam ist Anne Haring und Michael Mahren das Arbeiten in einem offenen und nicht in einem ergebnisorientierten Prozess. Die Interaktion mit den Gegebenheiten der Arbeitssituation, sich der Auseinandersetzung mit dem Material beziehungsweise mit den äußeren Einflüssen und mit dem Modell zu stellen, sie zu suchen ist für beide der künstlerische Impetus.

Maja Andrack, »über * oder über 1«

Maja Andrack (geb. Sokolova) zeigte in Schloss Dagstuhl von 14. Oktober bis 18. Dezember 2014 ihr neues künstlerisches Projekt über das Verschwinden. Das Konzept der 2014 entstandenen Arbeit „über * oder über 1"

Born in Baltersweiler, Germany, in 1938, Ohlmann studied at the Saarbrücken *Werkkunstschul* (now the Saarland Fine Arts Academy), where he later taught from 1967–1970, and at the Academy of Fine Arts in Stuttgart. Founder of the "Grouppe 7" artists' collective, Ohlmanm received several awards for his work during his lifetime, including the "Jerg Ratgeb"-Publikumspreis in 1979 and the Mia Münster Prize in 2003.

Press coverage of the exhibit described Aloys Ohlmann as a politically committed and internationally active artist, whose window paintings point to an emotional and spiritual experience that eludes direct interpretation.[77] Prof. Karl Otto Jung, a good friend of the artist during the latter's lifetime, introduced Ohlmann's work to exhibit opening guests. Each window image, he explained, "gives a view onto a different aspect of the world – sometimes cheerful, sometimes sad, sometimes scary, but always relevant to the world as it is today." Prof. Jung encouraged the guests to think of Ohlmann's words "Ich denke mehr, als in Wirklichkeit ist" ("I think more than what reality holds").

Anne Haring and Michael Mahren, »Sculptures and Drawings«

Schloss Dagstuhl hosted this dual-artist exhibit, featuring sculptural works by Anne Haring and drawings by Michael Mahren, from May to July 2014.

Anne Haring, a freelance photographer, sculptor, and artist, studied at the Staatliche Kunstakademie Düsseldorf under Prof. Karl Bobek and at the Cité Internationale des Arts in Paris. She later taught the drawing and photography Technical University of Darmstadt from 1999 to 2004.

Michael Mahren completed his artist studies at the Fachhochschule Saarbrücken with Oskar Holweck and Robert Sessler and now works as graphic designer. In 1980 he and Axel Reiß were awarded the first prize in the artistic design competition for the historic fire station of Saarbrücken, Germany.

Together, Haring and Mahren engage in an open, non-results-oriented process in which interacting with realities of the work situation and coming to terms with materials and external influences provide artistic impetus.

Maja Andrack, »über * oder über 1«

Maja Andrack (neé Sokolova) is a Saarland-based artist who uses different media, as determined by the context of each piece, to create works that also include narrative elements or texts. Maja Andrack's show at Schloss

[77] *Saarbrücker Zeitung* 24.03.2014, „Draußen vor dem Fenster: Nie gesehene Bilder von Aloys Ohlmann".

bestand aus einer Gegenüberstellung von S/W Fotografie und Malerei auf Leinwand. Maja Andrack arbeitet konzeptabhängig in verschiedenen Medien und bezieht narrative Elemente und Texte in ihren Arbeiten mit ein.

Dagstuhl, held from October 14 through December 18, 2014, presented her new artistic project on the concept of disappearance. Her 2014 work "über * oder über 1" insisted on a comparison between black and white photography and paintings on canvas.

Kunstankauf durch Spenden

10.2 *Art Sponsorship and Donations*

Schloss Dagstuhl durfte sich im Jahr 2014 über die Schenkung einer dreidimensionalen Reproduktion von M. C. Eschers sogenannter „Belvedere"-Zeichnung von Prof. Gershon Elber von Technion für seine permanente Kunstsammlung freuen (siehe Fig. 10.2). Prof. Elber ist ein Computer-Wissenschaftler, dessen Forschungsgebiet computergestützte, geometrische Konstruktion und Computergrafik einschließt. Mit Hilfe eines mehrlagigen Herstellungsverfahrens schuf Prof. Elber die außergewöhnliche, schöne Skulptur. Das Objekt steht momentan in einer Glasvitrine im Foyer des historischen Gebäudes von Schloss Dagstuhl.

Das Internetangebot von 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. Gleichzeitig ist es fruchtbar für das Zentrum, da die Möglichkeit besteht, Werke von Künstlern, die auf Schloss Dagstuhl ausgestellt haben, anzukaufen. Two works were purchased this way in 2014.

Weitere Spenden und Details werden in Kapitel 1 diskutiert. Nähere Informationen und aktuelle Neuigkeiten finden sich auf der Kunst-Webseite[78] von Dagstuhl. Wir möchten diese Stelle nutzen, allen Spendern, die 2014 zu der Kunstsammlung von Schloss Dagstuhl beigetragen haben, unseren Dank auszusprechen.

Schloss Dagstuhl was pleased to receive in 2014 as a gift for its permanent art collection a three-dimensional reproduction of M. C. Escher's so-called "impossible" Belvedere drawing from Prof. Gershon Elber of Technion (see Fig. 10.2). A computer scientist whose research interests include computer aided geometric designs and computer graphics, Prof. Elber created the unusual and beautiful sculpture with the help of a layered manufacturing process. The object currently stands in a glass vitrine in the foyer of the historic main building at Schloss Dagstuhl.

Dagstuhl's website contains a page featuring an Internet gallery enabling participants, individuals, and groups to make contributions to Dagstuhl for art donations. The works of art are featured 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. Two works were purchased this way in 2014.

More details about donations can be found in Chapter 1. For further information about the art exhibits at Schloss Dagstuhl and the art program in general, please visit Dagstuhl's art webpage[78]. We would like to take this opportunity to thank all of those who generously donated to the center's art collection in 2014.

[78] http://www.dagstuhl.de/art/

Fabian Treiber, »Neun Minuten vor Vegas« Works by artist Fabian Treiber | January 13 to March 07, 2014
Aloys Ohlmann (1938–2013), »Window Pictures – Hommage to Oskar Schlemmer« Memorial exhibition in honor of Aloys Ohlmann (†16.09.2013) Works by artist Aloys Ohlmann | March 10 to May 5, 2014
Anne Haring und Michael Mahren, »Sculptures and Drawings« Works by artists Anne Haring and Michael Mahren | May, 13 to July, 24, 2014
Maja Andrack, »über * oder über 1« Works by artist Maja Andrack | October 14 to December 18, 2014

Fig. 10.1
Art exhibitions in 2014.

Fig. 10.2
Sculptural reproduction of M. C. Escher's "impossible" Belvedere drawing. Donated to Schloss Dagstuhl by artist and computer scientist Prof. Gershon Elber in 2014.

11 Struktur der Gesellschaft
Structure of the Company

Gründung und Gesellschafter 11.1 *Formation and Shareholders*

Schloss Dagstuhl ist als eine gemeinnützige GmbH mit derzeit elf Gesellschaftern (siehe Fig. 11.1) organisiert. Dies sind die vier Gesellschafter, die Schloss Dagstuhl gegründet haben, nämlich die Gesellschaft für Informatik e.V. (GI), die Universität des Saarlandes, die Technische Universität Kaiserslautern und das Karlsruher Institut für Technologie. 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, Institut National de Recherche en Informatique et en Automatique (INRIA, Frankreich), Centrum Wiskunde & Informatica (CWI, Niederlande) und die Max-Planck-Gesellschaft (MPG, Deutschland) wurden 2005/2006 als weitere Gesellschafter aufgenommen.

Aufgrund eines Beschluss der Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung (BLK) 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 widmet sich mit herausragendem Erfolg seiner Aufgabe, die internationale Informatikforschung mit einem Seminarzentrum für wissenschaftliche Veranstaltungen zu unterstützen.

Im Mai 2014 trat ein geänderter Gesellschaftervertrag in Kraft. Die Änderungen beziehen sich im Wesentlichen auf das Wissenschaftliche Direktorium und werden in Abschnitt 11.3 dargestellt.

Schloss Dagstuhl is operated as a non-profit organization by currently eleven associates (cf. Fig. 11.1), including its four founding associates: the Gesellschaft für Informatik e.V.[79] (GI), the Universität des Saarlandes, the Technische Universität Kaiserslautern, and the Karlsruher Institut für Technologie. In 1994, the organization was extended to include four new associates: the Technischen 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, The Netherlands), and the Max-Planck-Gesellschaft (MPG, Germany).

By resolution of the “Bund-Länder-Kommission für Bildungsplanung und Forschungsförderung” (BLK),[80] the center has been classified as a research service institution for joint funding by the German federal and state governments since January of 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”[81] to “Schloss Dagstuhl – Leibniz-Zentrum für Informatik”[82] in 2008.

In July of 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.

A revised shareholders agreement came into effect in May, 2014. The amendments, which largely concern the Scientific Directorate, are outlined in Section 11.3.

Organe der Gesellschaft 11.2 *Dagstuhl Organs*

Die drei Organe von Schloss Dagstuhl – Leibniz-Zentrum für Informatik, die stellvertretend für die Gesellschaft als juristische Person handeln sind die folgenden:

- Gesellschafterversammlung
- Aufsichtsrat
- Geschäftsführung

Details zu den Organen sind den folgende Abschnitten zu entnehmen.

The three organs of Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 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.

[79] engl.: German Informatics Society
[80] engl.: Federal Government–State Commission for Educational Planning and Research Promotion
[81] engl.: International Conference and Research Center for Computer Science
[82] engl.: Schloss Dagstuhl – Leibniz Center for Informatics

Die Gesellschafterversammlung

Die Gesellschafter beschließen über alle Änderungen an der Gesellschaft, insbesondere über die Aufnahme weiterer Gesellschafter, über die Änderung des Gesellschaftsvertrags und über ihre Auflösung. Die Gesellschafter bestätigt 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. 11.2) setzen sich aus vier Repräsentanten der Gesellschaft für Informatik, je einem Vertreter der drei Gründungsuniversitäten, zwei Vertretern der später hinzugekommenen vier Universitäten und je einem Vertreter des Bundes und der beiden Bundesländer Saarland und Rheinland-Pfalz, in denen Schloss Dagstuhl formal den Sitz hat, zusammen. Die Amtszeit der Mitglieder des Aufsichtsrats 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 Wissenschaftliche 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 der Gründung von Schloss Dagstuhl in 1990 bis einschließlich April 2014 war Prof. Reinhard Wilhelm der Wissenschaftliche Direktor.

Shareholders' Meeting

All changes to the company, in particular the inclusion of new associates and the revision or even dissolution of the company agreement, 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 the once-yearly shareholders' meeting, 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 meaningful legal and economic manner. The board is involved in all essential matters with regard to research and financial planning.

The 12-member board (see Fig. 11.2) is composed of four representatives of the Gesellschaft für Informatik, one representative of 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 representatives of the federal and represented 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 of the Scientific Directorate, which he leads. Until April 2014 Prof. Reinhard Wilhelm was the Scientific Director of Schloss Dagstuhl, a position he had held since the founding of Schloss Dagstuhl in 1990. He was succeeded in May 2014 by Prof. Raimund Seidel, Ph.D..

Er wurde im Mai 2014 von Prof. Raimund Seidel, Ph. D. abgelöst.

Bis zur Neufassung des Gesellschaftsvertrags in 2014 wurde der Wissenschaftliche Direktor vom Wissenschaftlichen Direktorium aus dessen Mitte für eine Amtszeit von drei Jahren gewählt. Nach der Neufassung wird der Wissenschaftliche Direktor dem Aufsichtsrat von einer Findungskommission zur Berufung vorgeschlagen. Der Findungskommission gehören mindestens der Vorsitzende des Aufsichtsrats und der Vorsitzende des Wissenschaftlichen Beirats an. Die Amtszeit des Wissenschaftlichen Direktors beträgt fünf Jahre.

Die technischen und administrativen Aufgaben werden vom Technisch-administrativen Geschäftsführer wahrgenommen. Die Amtszeit des bisherigen Technisch-administrativen Geschäftsführers Dr. Christian Lindig endete im Juni 2014. Im Juli 2014 übernahm Frau Heike Meißner diese Position. Siehe Fig. 11.3.

Until the new shareholders' contract took effect, the Scientific Director was elected by the Scientific Directorate from their midst for a term of office of three years. With the legalization of the new contract, the Supervisory Board now appoints the Scientific Director on basis of the recommendation of a selection committee consisting of at least the chairperson of the Supervisory Board and the chairperson of the Scientific Advisory Board. The term of office of the Scientific Director is five years.

The Technical Administrative Director is responsible for technical and administrative tasks. The tenure of former Technical Administrative Director, Dr. Christian Lindig, ended in June 2014. In July 2015, he was succeeded by Heike Meißner. See Fig. 11.3.

Gremien der Gesellschaft 11.3 *Dagstuhl Bodies*

Die Organe von Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH werden durch drei Gremien unterstützt. Es sind die folgenden:

- Wissenschaftliches Direktorium
- Wissenschaftlicher Beirat
- Kuratorium

Details zu den Gremien werden in den folgenden Abschnitten ausgeführt.

Das Wissenschaftliche Direktorium

Das Wissenschaftliche Direktorium (siehe Fig. 11.4) ist für die Realisierung des Gesellschaftszwecks in fachlich-wissenschaftlicher Hinsicht verantwortlich. Es hat das Forschungs- und Veranstaltungsprogramm der Gesellschaft festzulegen, seine fachlich-wissenschaftliche Qualität zu sichern und seine Durchführung zu überwachen. Als wesentlicher Bestandteil dieser Aufgabe werden die Anträge auf Dagstuhl-Seminare und Dagstuhl-Perspektiven-Workshops von Mitgliedern des Wissenschaftlichen Direktoriums begutachtet. Auf den zweimal im Jahr stattfindenden Direktoriumssitzungen werden die Anträge diskutiert und es wird über ihre Annahme entschieden.

Die Art und Weise der Berufung der Mitglieder des Wissenschaftlichen Direktoriums hat sich mit dem Inkrafttreten des neuen Gesellschaftsvertrags erheblich geändert. Nach dem alten Gesellschaftsvertrag schlugen die Gesellschafter mindestens zwei Kandidaten für das Direktorium vor, die Gesellschaft für Informatik jedoch acht. Von diesen wählte der Wissenschaftliche Beirat je einen Kandidaten bzw. vier Kandidaten aus den Vorschlägen der Gesellschaft für Informatik aus und schlug sie dem Aufsichtsrat zur Berufung vor. Das Wissenschaftliche Direktorium bestand somit aus 14 Mitgliedern. Diese wählten aus ihrer Mitte den Wissenschaftlichen Direktor. Wurden nicht alle Teilgebiete der Informatik von den Mitgliedern des Direktoriums abge-

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. 11.4) is responsible for carrying out the company objectives from a technical and scientific point of view. It must determine the research and event program, ensure its technical and scientific quality, and monitor its execution. As a main task in support of this objective, members of the Scientific Directorate review proposals 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 procedure for appointing new members of the Scientific Directorate changed considerably with the new shareholders' contract. According to the old contract, each shareholder nominated at least two candidates for the board, with the exception of the Gesellschaft für Informatik which nominated eight candidates. Out of these, the Scientific Advisory Board selected for each shareholder one candidate and for the Gesellschaft für Informatik four. The Supervisory Board then appointed the members of the Scientific Directorate in accordance with this selection. The Directorate thus had 14 members and, furthermore, elected the Scientific Director from its midst. In the event that all relevant computer science topics were not covered by members of the Scientific Directorate, the Scientific Advisory Board nominated members-at-large for appoint-

deckt, konnte der Beirat *members-at-large* zur Berufung durch den Aufsichtsrat vorschlagen, die die fehlenden Teilgebiete für Gutachten zu Seminar- und Perspektiven-Workshop-Anträgen abdeckten (siehe Fig. 11.5).

Nach dem neuen Gesellschaftsvertrag wird nun die Position des Wissenschaftlichen Direktors unabhängig vom Wissenschaftlichen Direktorium besetzt, er gehört diesem aber an. Weiterhin empfiehlt nun der Wissenschaftliche Direktor dem Aufsichtsrat die Größe des Direktoriums. Neben den Gesellschaftern können nun auch das bestehende Wissenschaftliche Direktorium sowie auch der Beirat Kandidaten für das Wissenschaftliche Direktorium benennen. Die Auswahl der Kandidaten, die dem Aufsichtsrat zur Ernennung vorgeschlagen werden, obliegt nun dem Beirat zusammen mit dem Wissenschaftlichen Direktor. Die Bindung, dass eine feste Anzahl der Vorschläge eines Gesellschafters berücksichtigt werden muss, entfällt ersatzlos. Dieses Verfahren garantiert, dass die Gruppe möglicher Kandidaten zur Ernennung in das Direktorium nicht mehr eingeschränkt ist und dass alle relevanten Teilgebiete der Informatik durch Mitglieder des Direktoriums abgedeckt werden können. Die Members-at-large entfallen daher, die Amtszeit der bisherigen endete dementsprechend mit dem Inkrafttreten des geänderten Gesellschaftsvertrags in 2014.

Die Amtszeit der Mitglieder des Wissenschaftlichen Direktoriums außer der des Wissenschaftlichen Direktors beträgt nach wie vor 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 Entscheidungen zur Weiterentwicklung von Schloss Dagstuhl und bei der Gewinnung von Leitungspersonal. Zum anderen führt der Wissenschaftliche Beirat mindestens einmal zwischen je zwei Evaluierungen durch den Senatsausschuss Evaluierung der Leibniz-Gemeinschaft ein Audit durch, bei dem die gesamte Einrichtung begutachtet wird. Ein Bericht über das Audit wird der Leitung, dem Aufsichtsrat und dem Senatsausschuss vorgelegt.

Der Wissenschaftliche Beirat (siehe Fig. 11.6) 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.

ment by the Supervisory Board. These members-at-large were responsible for reviewing proposals on Dagstuhl Seminars and Dagstuhl Perspectives Workshops related to their topics (see Fig. 11.5).

The new shareholders contract stipulates that the Scientific Director should be appointed independently of the Scientific Directorate, although the Scientific Director still belongs to the Scientific Directorate. The Scientific Director 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 now done by the Scientific Advisory Board together with the Scientific Director. The constraint requiring participants in the selection process to consider a fixed number of candidates per shareholder was eliminated. This new procedure guarantees that the group of eligible candidates for appointment is no longer restricted, and that all relevant topical areas in computer science are covered by members of the Scientific Directorate. With this change, the role of members-at-large became superfluous and the position was dropped from the shareholders contract. The term of office for all present members-at-large ended with the coming into effect of the contract in 2014.

The term of office of Scientific Directorate members, with the exception of the Scientific Director, is still three years. It begins on November 1 of the year of appointment and ends three years later on October 31. A multiple reelection is 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 medium-term research and development plans and issues comments on the program budget draft, making recommendations on the use of resources. It also assists the Supervisory Board in important decisions with regard to future development and the acquisition of management staff and of the institute. 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. 11.6) 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 for a second term. The Scientific Advisory Board elects a chairperson from their midst. There is an annual meeting of the board. Members are appointed by the Supervisory Board in accordance with the suggestions of the Scientific Advisory Board.

Das Kuratorium

Das Kuratorium (siehe Fig. 11.7) 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 passende zugkräftige Organisatoren aus der Industrie zu gewinnen. Das Kuratorium wird regelmäßig aufgefordert, aus seinem Wirkungskreis passende Teilnehmer zu den Seminaren vorzuschlagen. Das industrielle Kuratorium tagt einmal im Jahr zusammen mit dem Wissenschaftlichen Beirat.

Industrial Advisory Board

The Industrial Advisory Board (see Fig. 11.7) performs a transmissional function between the center and the R&D departments and industry 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 five members whose term of office are five years. A reappointment for a second term is possible once. The members of the board 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 seminars known to it from its activities. It convenes once a year together with the Scientific Advisory Board.

Gesellschafter | Associates
Centrum Wiskunde & Informatica (CWI), The Netherlands
Gesellschaft für Informatik e.V., Germany
Institut National de Recherche en Informatique et en Automatique (INRIA), France
Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany
Karlsruher Institut für Technologie (KIT), Germany
Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V., Berlin, Germany
Technische Universität Darmstadt, Germany
Technische Universität Kaiserslautern, Germany
Universität des Saarlandes, Germany
Universität Stuttgart, Germany
Universität Trier, Germany

Fig. 11.1
Associates.

Aufsichtsrat | Supervisory Board
Dr. Doreen Becker Bundesministerium für Bildung und Forschung, Bonn, Germany | Representative of the German federal government
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.
Prof. Oliver Günther, Ph. D. Universität Potsdam, Germany | Representative of Gesellschaft für Informatik e.V. | *tenure ended in May 2014*
Wolfgang Habelitz Ministeriums für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of the Rhineland-Palatinate | *tenure ended in May 2014*
Prof. Dr.-Ing. Dr. h. c. Theo Härder Technische Universität Kaiserslautern, Germany | Representative of Technische Universität Kaiserslautern | *tenure ended in May 2014*
Dr. Christian Heimann Ministeriums für Bildung, Wissenschaft, Weiterbildung und Kultur, Mainz, Germany | Representative of Rhineland-Palatinate state | *tenure started in May, 2014*
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. | *tenure started in May 2014*
Prof. Dr. Volker Linneweber Universität des Saarlandes, 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 | *tenure started in May 2014*
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. 11.2
Supervisory Board members.

Geschäftsführung | Management
Christian Lindig (Technisch-administrative Geschäftsführer | Technical Administrative Director) Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH, Germany | *tenure ended in June 2014*
Heike Meißner (Technisch-administrative Geschäftsführerin | Technical Administrative Director) Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH, Germany | *tenure started in July 2014*
Prof. Raimund Seidel, Ph. D. (Wissenschaftlicher Direktor | Scientific Director) Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH and Universität des Saarlandes, Germany | *tenure started in May 2014*
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm (Wissenschaftlicher Direktor | Scientific Director) Schloss Dagstuhl – Leibniz Zentrum für Informatik GmbH and Universität des Saarlandes, Germany | *tenure ended in April 2014*

Fig. 11.3
Management.

Wissenschaftliches Direktorium | Scientific Directorate
Prof. Dr. Susanne Albers Humboldt-Universität, Berlin, Germany | *tenure ended in October 2014*
Prof. Dr. Bernd Becker Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Karsten Berns Technische Universität Kaiserslautern, Germany | *tenure ended in October 2014*
Prof. Dr. Stefan Diehl Universität Trier, Germany
Prof. Dr. Hans Hagen Technische Universität Kaiserslautern, Germany | *tenure started in November 2014*
Prof. Dr. Hannes Hartenstein Karlsruher Institut für Technologie, Germany
Prof. Dr.-Ing. Oliver Kohlbacher Eberhard Karls Universität Tübingen | *tenure started in November 2014*
Dr. Stephan Merz Institut National de Recherche en Informatique et en Automatique (INRIA), Nancy, France
Prof. Dr.-Ing. Bernhard Mitschang Universität Stuttgart, Germany
Prof. Dr. Bernhard Nebel Albert-Ludwigs-Universität Freiburg, Germany
Prof. Dr. Han La Poutré Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands | *tenure ended in October 2014*
Prof. Dr. Bernt Schiele Max-Planck-Institut für Informatik, Saarbrücken, Germany
Prof. Dr. Nicole Schweikardt Humboldt-Universität zu Berlin, Germany *(until August 2014 at Johann Wolfgang Goethe-Universität Frankfurt am Main, Germany)*
Prof. Raimund Seidel, Ph. D. Universität des Saarlandes, Germany
Prof. Dr. Michael Waidner Technische Universität Darmstadt, Germany
Prof. Dr.-Ing. Klaus Wehlre Rheinisch-Westfälische Technische Hochschule Aachen, Germany | *tenure started in November 2014*
Prof. Dr. Dr. h. c. Dr. h. c. Reinhard Wilhelm Universität des Saarlandes, Germany
Prof. Dr. Ir. Arjen P. de Vries Centrum Wiskunde & Informatica (CWI), Amsterdam, The Netherlands | *tenure started in November 2014*

Fig. 11.4
Scientific Directorate.

Members-at-Large (Board dissolved in May 2014)
Prof. Dr. Luca Benini ETH Zürich, Switzerland and University of Bologna, Italy | *tenure ended in May 2014*
Prof. em. Dr. Jan-Olof Eklundh Royal Institute of Technology, Stockholm, Sweden | *tenure ended in May 2014*
Prof. Dr. Friedemann Mattern ETH Zürich, Switzerland | *tenure ended in May 2014*

Fig. 11.5
Members-at-Large.

Wissenschaftlicher Beirat | Scientific Advisory Board
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), France
Prof. Dr. Mila E. Majster-Cederbaum Ludwig-Maximilians-Universität München, Germany | Chairwoman of the Scientific Advisory Board | *stepped down in September 2014*
Prof. Dr. Friedhelm Meyer auf der Heide Heinz Nixdorf Institut and Computer Science Departement, Universität Paderborn, Germany
Prof. Dr.-Ing. Dr. h. c. Andreas Reuter HITS GmbH, Heidelberg, Germany
Prof. em. Dr. Dr. h. c. Otto Spaniol RWTH Aachen, Germany
Dr. Anne Norekian Staatskanzlei des Saarlandes, Saarbrücken, Germany | *Guest*

Fig. 11.6
Scientific Advisory Board.

Kuratorium | Industrial Advisory Board
Dr. Udo Bub EIT ICT Labs, Berlin, Germany
Dr. Jorge R. Cuéllar Siemens AG, München, Germany
Dr.-Ing. Elmar Dorner SAP Research, Karlsruhe, Germany
Dr. Jo Ebergen Oracle Labs, United States
Dr. Goetz Graefe HP Labs, United States
Prof. Dr. Ralf Guido Herrtwich Daimler AG, Böblingen, Germany
Prof. Dr. Thomas Hofmann ETH Zürich, Switzerland *(until June 2013 at Google Research, Zürich, Switzerland)* | *tenure ended in May 2014*
Prof. Dr. Ulrich Lauther Siemens AG, München, Germany
Prof. Dr. Prabhakar Raghavan Google Inc. and Consulting Professor at Stanford University, United States
Prof. Dr.-Ing. Dr. h. c. Andreas Reuter HITS GmbH, Heidelberg, Germany
Prof. Dr. Volker Tresp Siemens AG, München, Germany and Ludwig-Maximilians-Universität München, Germany

Fig. 11.7
Industrial Advisory Board.

12 Statistik 2014
Statistics 2014

Statistiken zum wissenschaftlichen Programm

12.1

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. 12.1 dargestellt.

Veranstaltungs-bezogene Daten: Fig. 12.2 zeigt Anzahl und Anteil der eingeladenene Seminarteilnehmer, welche die Einladung annehmen bzw. ablehnen. Daten zu Anzahl, Größe und Dauer der durchgeführten Seminare sind in Fig. 12.3 angegeben. Fig. 12.4 zeigt die Verteilung in Bezug auf die verschiedenen Veranstaltungstypen.

Teilnehmer-bezogene Daten: Die Teilnehmerzahlen — abhängig vom Veranstaltungstyp — gibt Fig. 12.5 an. Ein Überblick über die Ergebnisse der reglmäßigen Gästebefragungen kann 12.6 entnommen werden. Fig. 12.7 zeigt die Verteilung der Herkunfsländer unserer Gäste. Die Anzahl von früheren Seminarbesuchen kann man Fig. 12.8 entnehmen. Fig. 12.9 gibt Auskunft über die Altersstruktur der Teilnehmer. Schlussendlich findet man in Fig. 12.10 Zahlen zu den Gasttagen.

Gender-bezogene Daten: Fig. 12.11 enthält Daten zur Gender-Komposition der Seminarleitung. Die Abbildungen Fig. 12.12 und Fig. 12.13 zeigen insbesondere die Anteile weiblicher Teilnehmer bzw. Einladungen an weibliche Wissenschaftler.

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. 12.1 shows the number of submitted proposals for Dagstuhl Seminars and Dagstuhl Perspectives Workshops, as well as acceptance rates for recent years.

Event-related data: Fig. 12.2 shows the number and the fraction of invited seminar participants who accepted and declined the invitation. Fig. 12.5 provides data about the number, size, and duration of the seminars held at Dagstuhl. Fig. 12.4 shows the distribution of different types of events at Dagstuhl.

Participant-related data: Fig. 12.5 shows the number of participants according to event type. An overview of the results of the participant survey for Dagstuhl Seminars and Dagstuhl Perspectives Workshops can be found in 12.6. Fig. 12.7 shows the distribution of country affiliations. Fig. 12.8 displays how often participants have attended seminars in the past. Fig. 12.9 gives data on the seniority of participants. Finally, Fig. 12.10 states the number of guest days hosted at Schloss Dagstuhl.

Gender-related data: Fig. 12.11 shows mixed-gender data. In Fig. 12.12 and Fig. 12.13 data is given with regard to female participants and invitees, respectively.

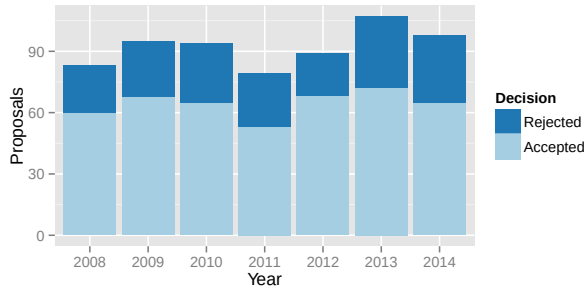

(a)Graphical distribution for 2008–2014

Year	Proposals	Accepted		Rejected	
	#	#	%	#	%
2008	83	60	72.3	23	27.7
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

(b)Detailed numbers for 2008–2014

Fig. 12.1
Proposals and acceptance rates for Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

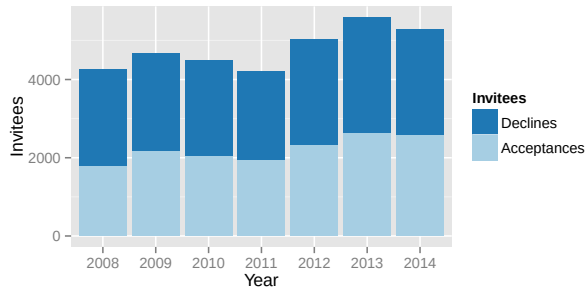

(a)Graphical distribution for 2008–2014

Year	Invitees	Accepances		Declines	
	#	#	%	#	%
2008	4268	1801	42.2	2467	57.8
2009	4670	2168	46.4	2502	53.6
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

(b)Detailed numbers for 2008–2014

Fig. 12.2
Invitees, accepting invitees, and declining invitees to Dagstuhl Seminars and Dagstuhl Perspectives Workshops.

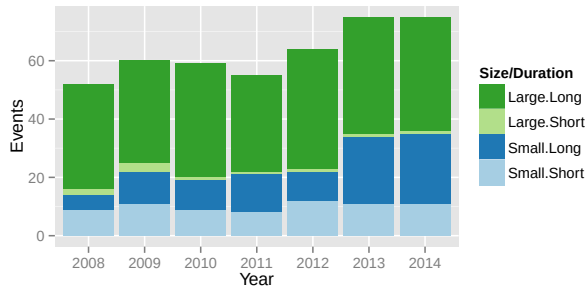

(a)Graphical distribution for 2008–2014

Year	30-person seminars		45-person seminars		Total
	3-day	5-day	3-day	5-day	
2008	9	5	2	36	52
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

(b)Detailed numbers for 2008–2014

Fig. 12.3
Size and duration of accepted Dagstuhl Seminars and Dagstuhl Perspectives Workshops. Small = 30-person seminar, Large = 45-person seminar, Short = 3-day seminar, Long = 5-day seminar.

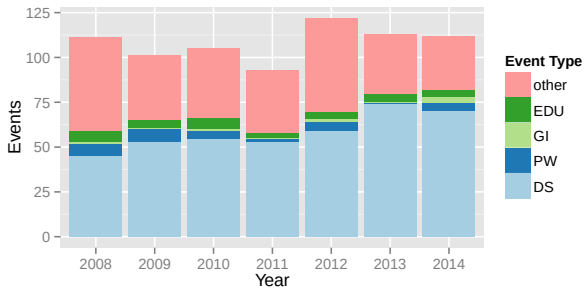

(a)Graphical distribution for 2008–2014

Year	DS	PW	GI	EDU	OE	Total
2008	45	7	1	6	52	111
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

(b)Detailed numbers for 2008–2014

Fig. 12.4
Number of all events held at Dagstuhl, by type. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, OE = other event.

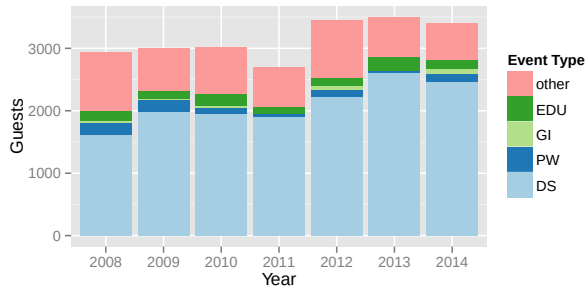

(a)Graphical distribution for 2008–2014

Year	DS		PW		GI		EDU		OE		Total
	#	%	#	%	#	%	#	%	#	%	#
2008	1622	55.3	179	6.1	32	1.1	166	5.7	932	31.8	2931
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

(b)Detailed numbers for 2008–2014

Fig. 12.5
Number of participants. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, OE = other event.

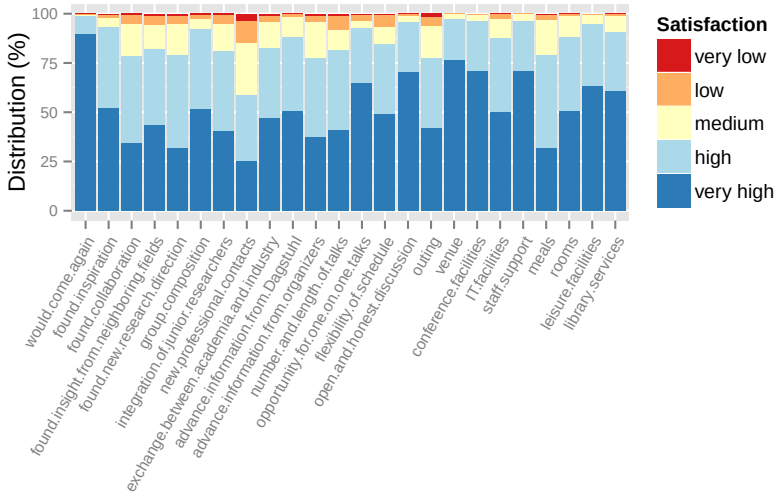

(a)Graphical distribution for 2014

	2008	2009	2010	2011	2012	2013	2014	2014 – 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	2	4	11	127	1286	1430
found inspiration	4.4	4.3	4.3	4.4	4.4	4.5	4.4	5	25	68	591	754	1443
found collaboration	3.8	3.9	3.9	4.0	4.0	4.0	4.1	9	65	232	630	493	1429
found insight from neighboring fields	4.0	4.1	4.0	4.1	4.1	4.1	4.2	14	66	173	561	622	1436
found new research direction	3.8	3.9	3.9	4.0	4.0	4.0	4.0	12	61	219	673	447	1412
group composition	4.3	4.4	4.4	4.4	4.4	4.4	4.4	9	26	76	588	745	1444
integration of junior researchers	4.2	4.2	4.2	4.2	4.2	4.2	4.2	11	64	190	580	580	1425
new professional contacts	3.5	3.5	3.5	3.7	3.6	3.7	3.7	48	161	365	470	359	1403
exchange between academia and industry	4.1	4.1	4.3	4.1	4.3	4.2	4.2	8	27	113	309	406	863
advance information from Dagstuhl	4.2	4.2	4.3	4.3	4.4	4.4	4.4	1	21	141	540	717	1420
advance information from organizers	4.0	4.1	4.0	4.2	4.1	4.1	4.1	14	45	252	558	523	1392
number and length of talks	4.0	4.0	4.1	4.1	4.1	4.2	4.1	12	106	143	589	588	1438
opportunity for one on one talks	4.5	4.4	4.5	4.5	4.4	4.5	4.5	6	49	50	396	933	1434
flexibility of schedule	4.2	4.2	4.3	4.2	4.2	4.2	4.3	7	85	127	502	707	1428
open and honest discussion	4.7	4.7	4.7	4.7	4.6	4.7	4.7	4	12	38	367	1015	1436
outing	4.2	4.1	4.0	4.2	4.1	4.1	4.1	18	44	178	386	450	1076
venue	4.7	4.7	4.7	4.7	4.7	4.7	4.7	0	3	33	300	1099	1435
conference facilities	4.7	4.8	4.7	4.8	4.7	4.6	4.7	0	7	48	359	1015	1429
IT facilities	4.5	4.5	4.5	4.6	4.4	4.4	4.4	2	28	128	483	653	1294
staff support	4.7	4.7	4.7	4.7	4.7	4.7	4.7	0	4	43	350	963	1360
meals	4.3	4.3	4.2	4.3	4.2	4.1	4.1	7	37	254	674	459	1431
rooms	4.5	4.5	4.5	4.5	4.4	4.4	4.4	3	9	154	541	721	1428
leisure facilities	4.6	4.5	4.5	4.6	4.6	4.6	4.6	0	4	64	427	850	1345
library services	4.7	4.6	4.6	4.5	4.5	4.5	4.5	1	5	62	215	439	722

(b)Averages for 2008–2014 and detailed numbers for 2014: 1 = very low, 2 = low, 3 = medium, 4 = high, 5 = very high

Fig. 12.6
Satisfaction of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants. According to survey results.

Country	A	B	Total
Germany	614	595	1209
United States	509	23	532
United Kingdom	245	22	267
France	204	26	230
Netherlands	129	8	137
Italy	96	5	101
Switzerland	80	10	90
Canada	81	6	87
Israel	62	2	64
Austria	43	13	56
Luxembourg	14	41	55
Sweden	43	12	55
Denmark	46	6	52
Belgium	46	3	49
Norway	35	5	40
Australia	27	10	37
Japan	35	0	35
Spain	30	3	33
Czech Republic	27	3	30
Finland	25	3	28
Poland	22	2	24
Ireland	20	2	22
India	21	0	21
Brazil	13	4	17
China	14	0	14
Russian Federation	14	0	14
Greece	8	3	11
Portugal	8	1	9
Hungary	8	0	8
Republic of Korea	7	1	8
Estonia	6	1	7
Singapore	7	0	7
Chile	5	0	5
Hong Kong	5	0	5
Romania	4	1	5
Slovenia	5	0	5
Iceland	4	0	4
New Zealand	1	3	4
South Africa	4	0	4
Malta	3	0	3
Qatar	3	0	3
Saudi Arabia	3	0	3
Turkey	3	0	3
Argentina	2	0	2
Serbia	2	0	2
Tunisia	0	2	2
Bulgaria	0	1	1
Colombia	0	1	1
Croatia	0	1	1
Egypt	1	0	1
Lithuania	1	0	1
Mexico	1	0	1
Taiwan	1	0	1
Trinidad and Tobago	1	0	1
United Arab Emirates	1	0	1
Venezuela	1	0	1
Total	2590	819	3409

(a)Details for 2014 by country

Region	A		B		Total	
	#	%	#	%	#	%
Europe (w/o Germany)	1171	45.2	172	21	1343	39.4
Germany	614	23.7	595	72.6	1209	35.5
North America	590	22.8	29	3.5	619	18.2
Asia	159	6.1	3	0.4	162	4.8
Australia	28	1.1	13	1.6	41	1.2
South America	23	0.9	5	0.6	28	0.8
Africa	5	0.2	2	0.2	7	0.2
Total	2590	100	819	100	3409	100

(b)Details for 2014 by region

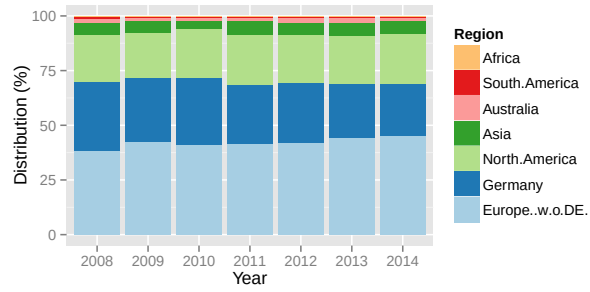

(c)Graphical distribution of seminar type A in 2008–2014 by year and region

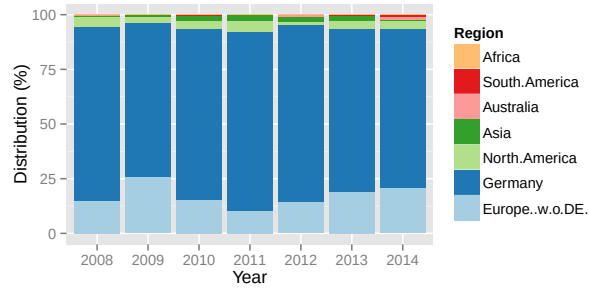

(d)Graphical distribution of seminar type B in 2008–2014 by year and region

Fig. 12.7

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 other events).

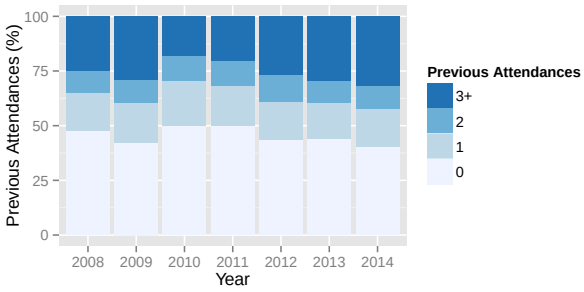

(a)Graphical distribution for 2008–2014

Year	Number of Previous Attendances								Total
	0		1		2		>2		
	#	%	#	%	#	%	#	%	#
2008	257	48	92	17	56	10	133	25	538
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

(b)Detailed numbers for 2008–2014

Fig. 12.8

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.

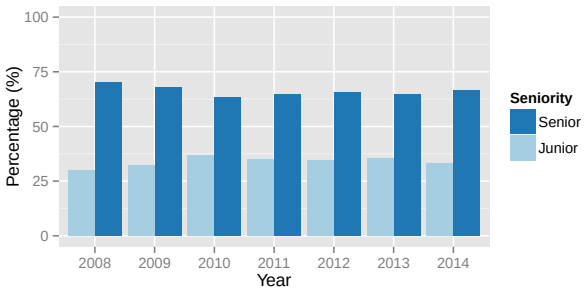

(a)Graphical distribution for 2008–2014

Year	Junior		Senior		Total
	#	%	#	%	#
2008	131	30.0	305	70.0	436
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

(b)Detailed numbers for 2008–2014

Fig. 12.9

Self-assigned seniority of Dagstuhl Seminar and Dagstuhl Perspectives Workshop participants.

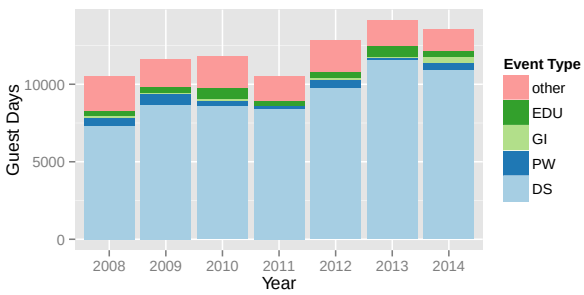

(a)Graphical distribution for 2008–2014

Year	DS	PW	GI	EDU	OE	Total
2008	7309	525	109	379	2226	10548
2009	8717	657	77	378	1776	11605
2010	8572	381	125	722	2002	11802
2011	8415	228	0	266	1604	10513
2012	9798	458	190	393	2031	12870
2013	11612	130	0	753	1614	14109
2014	10939	475	348	390	1370	13522

(b)Detailed numbers for 2008–2014

Fig. 12.10

Number of guest days at Schloss Dagstuhl. DS = Dagstuhl Seminar, PW = Dagstuhl Perspectives Workshop, GI = GI-Dagstuhl Seminar, EDU = educational event, OE = other event.

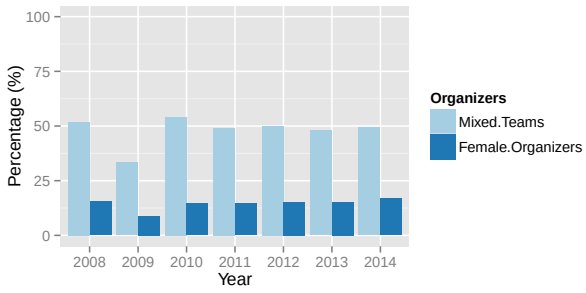

(a)Graphical distribution for 2008–2014

	Organizer Teams			Organizers		
Year	Total	Mixed		Total	Female	
	#	#	%	#	#	%
2008	52	27	51.9	200	31	15.5
2009	60	20	33.3	228	20	8.8
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

(b)Detailed numbers for 2008–2014

Fig. 12.11
Dagstuhl Seminars and Dagstuhl Perspectives Workshops with mixed-gender organizer teams.

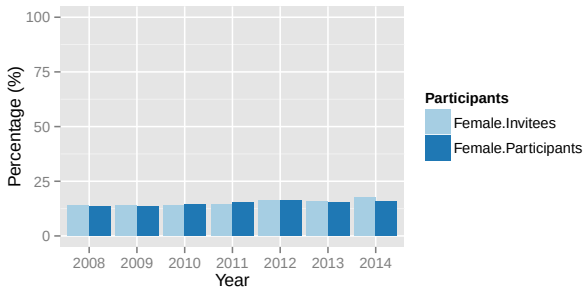

(a)Graphical distribution for 2008–2014

	Invitees			Participants		
Year	Total	Female		Total	Female	
	#	#	%	#	#	%
2008	4268	595	13.9	1801	244	13.5
2009	4670	651	13.9	2168	296	13.7
2010	4499	630	14.0	2053	293	14.3
2011	4223	602	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

(b)Detailed numbers for 2008–2014

Fig. 12.12
Female invitees and participants in Dagstuhl Seminars and Dagstuhl Perspectives Workshops, by year.

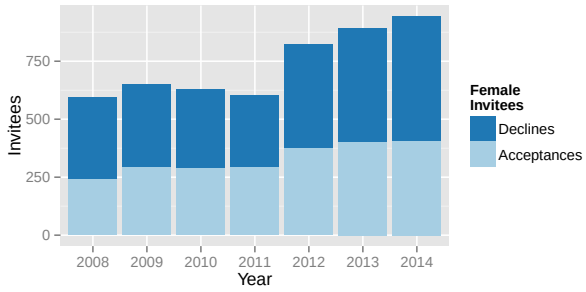

(a)Graphical distribution for 2008–2014

Year	Female Invitees	Acceptances		Declines	
	#	#	%	#	%
2008	595	244	41.0	351	59.0
2009	651	296	45.5	355	54.5
2010	630	293	46.5	337	53.5
2011	602	295	49.0	307	51.0
2012	822	377	45.9	445	54.1
2013	889	401	45.1	488	54.9
2014	943	406	43.1	537	56.9

(b)Detailed numbers for 2008–2014

Fig. 12.13
Female invitees to Dagstuhl Seminar and Dagstuhl Perspectives Workshops.

12.2 Statistiken zur Bibliographiedatenbank dblp

Dieser Abschnitt enthält statistische Daten zur Bibliographiedatenbank dblp. Fig. 12.14 listet die durchschnittlichen Nutzungszahlen in 2014. Ein Überblick über die Entwicklung des dblp Datenbestandes kann Fig. 12.15 und Fig. 12.16 entnommen werden.

12.2 *Statistics of the dblp computer science bibliography*

This section provides statistical data about the dblp computer science bibliography. Fig. 12.14 show the average usage statistics of the dblp servers in 2014. An overview of the development of the dblp database can be found in Fig. 12.15 and Fig. 12.16.

	Trier 1	Trier 2	Dagstuhl
user sessions (visits) per day	21,057	4,703	326
page views per day	174,247	47,531	14,964
page views per user session	8.2	10.1	45.8
distinct users (IPs) per month	327,299	76,566	4,399
data served per month	825.2 GB	345.8 GB	27.2 GB
as above, including bots	5,187.5 GB	672.2 GB	56.4 GB

Fig. 12.14
Average usage of the three dblp servers in 2014. Trier 1 = http://informatik.uni-trier.de/~ley/db, Trier 2 = http://dblp.uni-trier.de, Dagstuhl = http://dblp.dagstuhl.de

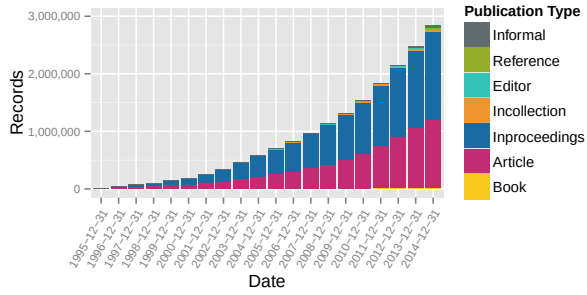

(a)Graphical distribution for 1996–2014

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2008	1,258	0.1	423,844	37.5	684,105	60.6	8,323	0.7	11,460	1.0	0	0.0	8	0.0	1,128,998
2009	1,350	0.1	508,111	38.5	781,394	59.3	14,507	1.1	13,006	1.0	0	0.0	8	0.0	1,318,376
2010	1,435	0.1	606,110	39.7	890,669	58.3	15,098	1.0	14,501	0.9	0	0.0	8	0.0	1,527,821
2011	9,351	0.5	744,807	40.6	1,042,917	56.9	19,636	1.1	16,889	0.9	0	0.0	13	0.0	1,833,613
2012	16,032	0.7	891,026	41.3	1,195,321	55.4	21,771	1.0	19,898	0.9	961	0.0	13,298	0.6	2,158,307
2013	16,827	0.7	1,042,792	42.0	1,346,764	54.2	24,951	1.0	22,539	0.9	961	0.0	29,165	1.2	2,483,999
2014	17,546	0.6	1,174,381	41.4	1,544,660	54.4	27,081	1.0	25,960	0.9	1,992	0.1	47,782	1.7	2,839,402

(b)Detailed numbers for 2008–2014

Fig. 12.15
Development of the total size of the dblp database.

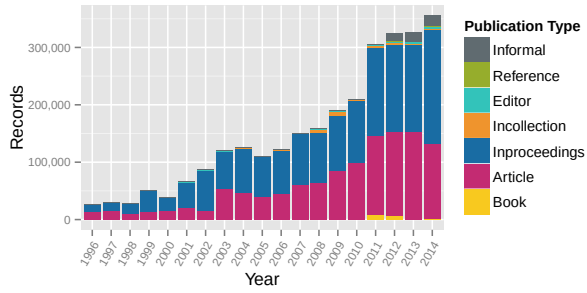

(a)Graphical distribution for 1996–2014

Year	Book		Article		Inproceedings		Incollection		Editor		Reference		Informal		Total
	#	%	#	%	#	%	#	%	#	%	#	%	#	%	#
2008	45	0.0	63,277	39.7	88,471	55.5	5,731	3.6	1,842	1.2	0	0.0	0	0.0	159,366
2009	92	0.0	84,267	44.5	97,289	51.4	6,184	3.3	1,546	0.8	0	0.0	0	0.0	189,378
2010	85	0.0	97,999	46.8	109,275	52.2	591	0.3	1,495	0.7	0	0.0	0	0.0	209,445
2011	7,916	2.6	138,697	45.4	152,248	49.8	4,538	1.5	2,388	0.8	0	0.0	5	0.0	305,792
2012	6,681	2.1	146,219	45.0	152,404	46.9	2,135	0.7	3,009	0.9	961	0.3	13,285	4.1	324,694
2013	795	0.2	151,766	46.6	151,443	46.5	3,180	1.0	2,641	0.8	0	0.0	15,867	4.9	325,692
2014	719	0.2	131,589	37.0	197,896	55.7	2,130	0.6	3,421	1.0	1,031	0.3	18,617	5.2	355,403

(b)Detailed numbers for 2008–2014

Fig. 12.16
Rate of inclusion of new publications to dblp.

13 Veranstaltungen 2014
Schedule of Events 2014

Dagstuhl-Seminare 13.1 *Dagstuhl Seminars*

14021 – Symmetric Cryptography
Frederik Armknecht (Universität Mannheim, DE), Helena Handschuh (Cryptography Research Inc. – San Francisco, US), Tetsu Iwata (Nagoya University, JP), Bart Preneel (KU Leuven, BE)
January 5–10, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14021

14031 – Randomized Timed and Hybrid Models for Critical Infrastructures
Erika Abraham (RWTH Aachen, DE), Alberto Avritzer (Siemens – Princeton, US), Anne Remke (University of Twente, NL), William H. Sanders (University of Illinois – Urbana Champaign, US)
January 12–17, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14031

14032 – Planning with epistemic goals
Thomas Agotnes (University of Bergen, NO), Gerhard Lakemeyer (RWTH Aachen, DE), Benedikt Löwe (University of Amsterdam, NL), Bernhard Nebel (Universität Freiburg, DE)
January 12–15, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14032

14041 – Quantitative Models: Expressiveness, Analysis, and New Applications
Manfred Droste (Universität Leipzig, DE), Paul Gastin (ENS – Cachan, FR), Kim Guldstrand Larsen (Aalborg University, DK), Axel Legay (INRIA Rennes – Bretagne Atlantique, FR)
January 19–24, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14041

14042 – Do-it-yourself Networking: an Interdisciplinary Approach
Panayotis Antoniadis (ETH Zürich, CH), Jörg Ott (Aalto University, FI), Andrea Passarella (CNR – Pisa, IT)
January 19–22, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14042

14051 – Algorithms for Wireless Communication
Guy Even (Tel Aviv University, IL), Magnus M. Halldorsson (Reykjavik University, IS), Yvonne-Anne Pignolet (ABB Corporate Research – Baden-Dättwil, CH), Christian Scheideler (Universität Paderborn, DE)
January 26–31, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14051

14052 – Ethics in Data Sharing
Julie E. Cohen (Georgetown Univ. – Washington, US), Sven Dietrich (Stevens Institute of Technology, US), Mireille Hildebrandt (Free University of Brussels, BE), Aiko Pras (University of Twente, NL), Lenore D. Zuck (University of Illinois – Chicago, US)
January 26–31, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14052

14061 – Statistical Techniques for Translating to Morphologically Rich Languages
Alexander M. Fraser (LMU München, DE), Kevin Knight (USC – Marina del Rey, US), Philipp Koehn (University of Edinburgh, GB), Helmut Schmid (LMU München, DE), Hans Uszkoreit (Universität des Saarlandes, DE)
February 2–7, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14061

14062 – The Pacemaker Challenge: Developing Certifiable Medical Devices
Dominique Méry (LORIA – Nancy, FR), Daniel Ratiu (fortiss GmbH – München, DE), Bernhard Schätz (fortiss GmbH – München, DE), Alan Wassyng (McMaster University – Hamilton, CA)
February 2–7, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14062

14071 – Graph Modification Problems
Hans L. Bodlaender (Utrecht University and Technical University Eindhoven), Pinar Heggernes (University of Bergen, NO), Daniel Lokshtanov (University of Bergen, NO)
February 9–14, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14071

14072 – New Perspectives in Shape Analysis
Michael Breuß (BTU Cottbus, DE), Alfred M. Bruckstein (Technion – Haifa, IL), Petros Maragos (National TU Athens, GR), Stefanie Wuhrer (Universität des Saarlandes, DE)
February 9–14, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14072

14081 – Robots Learning from Experiences
Anthony G. Cohn (University of Leeds, GB), Bernd Neumann (Universität Hamburg, DE), Alessandro Saffiotti (University of Örebro, SE), Markus Vincze (TU Wien, AT)
February 16–21, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14081

14082 – Visualization and Processing of Higher Order Descriptors for Multi-Valued Data
Bernhard Burgeth (Universität des Saarlandes, DE), Ingrid Hotz (DLR – Braunschweig, DE), Anna Vilanova Bartroli (TU Delft, NL), Carl-Fredrik Westin (Harvard Medical School – Boston, US)
February 16–21, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14082

14091 – Data Structures and Advanced Models of Computation on Big Data
Alejandro Lopez-Ortiz (University of Waterloo, CA), Ulrich Carsten Meyer (Goethe-Universität Frankfurt am Main, DE), Robert Sedgewick (Princeton University, US)
February 23–28, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14091

14092 – Digital Evidence and Forensic Readiness
Glenn S. Dardick (Longwood University – Farmville, US), Barbara Endicott-Popovsky (University of Washington – Seattle, US), Pavel Gladyshev (University College Dublin, IE), Thomas Kemmerich (Gjøvik University College, NO), Carsten Rudolph (Fraunhofer SIT – Darmstadt, DE)
February 23–28, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14092

14101 – Preference Learning
Johannes Fürnkranz (TU Darmstadt, DE), Eyke Hüllermeier (Universität Marburg, DE), Cynthia Rudin (MIT, US), Scott Sanner (NICTA – Canberra, AU), Roman Slowinski (Poznan University of Technology, PL)
March 2–7, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14101

14111 – Combinatorics and Algorithmics of Strings
Maxime Crochemore (King's College London, GB), James D. Currie (University of Winnipeg, CA), Gregory Kucherov (University Paris-Est – Marne-la-Vallée, FR), Dirk Nowotka (Universität Kiel, DE)
March 9–14, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14111

14121 – Computational Complexity of Discrete Problems
Anna Gál (University of Texas – Austin, US), Michal Koucký (Charles University – Prague, CZ), Oded Regev (New York University, US), Rüdiger Reischuk (Universität Lübeck, DE)
March 16–21, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14121

14122 – Verification of Cyber-Physical Systems
Rupak Majumdar (MPI-SWS – Kaiserslautern, DE), Richard M. Murray (CalTech – Pasadena, US), Pavithra Prabhakar (IMDEA Software – Madrid, ES)
March 16–21, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14122

14131 – Computational Models of Cultural Behavior for Human-Agent Interaction
Elisabeth André (Universität Augsburg, DE), Ruth Aylett (Heriot-Watt University Edinburgh, GB), Gert Jan Hofstede (Wageningen University, NL), Ana Paiva (INESC-ID – Porto Salvo, PT)
March 23–28, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14131

14132 – Interaction and Collective Movement Processing
Maike Buchin (Ruhr-Universität Bochum, DE), Luca Giuggioli (University of Bristol, GB), Guy Theraulaz (Université Paul Sabatier – Toulouse, FR), Marc van Kreveld (Utrecht University, NL)
March 23–28, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14132

14141 – Reachability Problems for Infinite-State Systems
Javier Esparza (TU München, DE), Alain Finkel (ENS – Cachan, FR), Pierre McKenzie (University of Montréal, CA), Joel Ouaknine (University of Oxford, GB)
March 30 to April 4, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14141

14142 – Spatial reference in the Semantic Web and in Robotics
Aldo Gangemi (CNR – Rome, IT), Verena V. Hafner (HU Berlin, DE), Werner Kuhn (University of California – Santa Barbara, US), Simon Scheider (Universität Münster, DE), Luc Steels (Free University of Brussels, BE)
March 30 to April 4, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14142

14171 – Evaluating Software Verification Systems: Benchmarks and Competitions
Dirk Beyer (Universität Passau, DE), Marieke Huisman (University of Twente, NL), Vladimir Klebanov (KIT – Karlsruher Institut für Technologie, DE), Rosemary Monahan (NUI Maynooth, IE)
April 21–25, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14171

14172 – Unifying Product and Software Configuration
Krzysztof Czarnecki (University of Waterloo, CA), Arnaud Hubaux (ASML – Veldhoven, NL), Ethan Jackson (Microsoft Corp. – Redmond, US), Dietmar Jannach (TU Dortmund, DE), Tomi Männistö (University of Helsinki, FI)
April 21–24, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14172

14181 – Multi-agent systems and their role in future energy grids
Michael N. Huhns (University of South Carolina – Columbia, US), Wolfgang Ketter (Erasmus University – Rotterdam, NL), Ryszard Kowalczyk (Swinburne University of Technology – Melbourne, AU), Fabrice Saffre (BT Research – Ipswich, GB), Rainer Unland (Universität Duisburg-Essen, DE)
April 27 to May 2, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14181

14201 – Horn formulas, directed hypergraphs, lattices and closure systems: related formalisms and applications
Kira V. Adaricheva (Yeshiva University – New York, US), Giuseppe F. Italiano (University of Rome "Tor Vergata", IT), Hans Kleine Büning (Universität Paderborn, DE), György Turan (University of Illinois – Chicago, US)
May 11–16, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14201

14202 – JA4AI – Judgment Aggregation for Artificial Intelligence
Franz Dietrich (CNRS, France & University of East Anglia, GB), Ulle Endriss (University of Amsterdam, NL), Davide Grossi (University of Liverpool, GB), Gabriella Pigozzi (University Paris-Dauphine, FR), Marija Slavkovik (University of Bergen, NO)
May 12–15, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14202

14211 – The Future of Refactoring
Danny Dig (Oregon State University, US), William G. Griswold (University of California – San Diego, US), Emerson Murphy-Hill (North Carolina State University, US), Max Schaefer (Semmle Ltd. – Oxford, GB)
May 18–23, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14211

14221 – Geometric Modeling
Falai Chen (Univ. of Science & Technology of China – Anhui, CN), Tor Dokken (SINTEF IKT Applied Mathematics – Oslo, NO), Thomas A. Grandine (The Boeing Company – Seattle, US), Stefanie Hahmann (University of Grenoble/INRIA, LJK-CNRS, FR)
May 25–30, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14221

14231 – Scientific Visualization
Min Chen (University of Oxford, GB), Charles D. Hansen (University of Utah, US), Penny Rheingans (University of Maryland, Baltimore Country, US), Gerik Scheuermann (Universität Leipzig, DE)
June 1–6, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14231

14232 – Design and Synthesis from Components
Jakob Rehof (TU Dortmund, DE), Moshe Y. Vardi (Rice University, US)
June 1–6, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14232

14241 – Challenges in Analysing Executables: Scalability, Self-Modifying Code and Synergy
Roberto Giacobazzi (University of Verona, IT), Axel Simon (TU München, DE), Sarah Zennou (Airbus Group – Suresnes, FR)
June 9–13, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14241

14261 – Software Development Analytics
Harald Gall (Universität Zürich, CH), Tim Menzies (West Virginia University – Morgantown, US), Laurie Williams (North Carolina State University, US), Thomas Zimmermann (Microsoft Corp. – Redmond, US)
June 22–27, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14261

14271 – Scripting Languages and Frameworks: Analysis and Verification
Fritz Henglein (University of Copenhagen, DK), Ranjit Jhala (University of California – San Diego, US), Shriram Krishnamurthi (Brown University, US), Peter Thiemann (Universität Freiburg, DE)
June 29 to July 4, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14271

14281 – Feature Interactions: The Next Generation
Sven Apel (Universität Passau, DE), Joanne M. Atlee (University of Waterloo, CA), Luciano Baresi (Politecnico di Milano University, IT), Pamela Zave (AT&T Labs Research – Bedminster, US)
July 6–11, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14281

14282 – Crowdsourcing and the Semantic Web
Abraham Bernstein (Universität Zürich, CH), Jan Marco Leimeister (Universität Kassel, DE & Universität St. Gallen, CH), Natasha Noy (Google Inc. – Mountain View, US), Cristina Sarasua (Universität Koblenz-Landau, DE), Elena Simperl (University of Southampton, GB)
July 6–9, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14282

14291 – Information-Centric Networking 3
Dirk Kutscher (NEC Laboratories Europe – Heidelberg, DE), Taekyoung Kwon (Seoul National University, KR), Ignacio Solis (Xerox PARC – Palo Alto, US)
July 13–16, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14291

14292 – Network Attack Detection and Defense: Securing Industrial Control Systems for Critical Infrastructures
Marc C. Dacier (QCRI – Doha, QA), Frank Kargl (Universität Ulm, DE), Hartmut König (BTU Cottbus, DE), Alfonso Valdes (University of Illinois – Urbana Champaign, US)
July 13–16, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14292

14301 – Computational Humanities – bridging the gap between Computer Science and Digital Humanities
Chris Biemann (TU Darmstadt, DE), Gregory R. Crane (Tufts University, US), Christiane D. Fellbaum (Princeton University, US), Alexander Mehler (Goethe-Universität Frankfurt am Main, DE)
July 20–25, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14301

14302 – Digital Palaeography: New Machines and Old Texts
Tal Hassner (The Open University of Israel – Raanana, IL), Robert Sablatnig (TU Wien, AT), Dominique Stutzmann (CNRS – Paris, FR), Ségolène Tarte (University of Oxford, GB)
July 20–24, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14302

14331 – Querying and Reasoning Under Expressive Constraints
Michael Benedikt (University of Oxford, GB), Carsten Lutz (Universität Bremen, DE), Balder Ten Cate (University of California – Santa Cruz, US)
August 10–14, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14331

14332 – Formal Methods for Coordinating Multi-Agent Systems
Thomas Agotnes (University of Bergen, NO), Nils Bulling (TU Clausthal, DE), Sascha Ossowski (University Rey Juan Carlos, ES)
August 10–14, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14332

14341 – Resource-bounded Problem Solving
Yll Haxhimusa (TU Wien, AT), Iris van Rooij (Radboud University Nijmegen, NL), Sashank Varma (University of Minnesota – Minneapolis, US), Todd Wareham (Memorial University of Newfoundland, CA)
August 17–22, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14341

14342 – Equilibrium Computation
Nimrod Megiddo (IBM Almaden Center, US), Kurt Mehlhorn (MPI für Informatik – Saarbrücken, DE), Rahul Savani (University of Liverpool, GB), Vijay V. Vazirani (Georgia Institute of Technology, US), Mihalis Yannakakis (Columbia University – New York, US)
August 17–22, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14342

14351 – Decision Procedures and Abstract Interpretation
Daniel Kroening (University of Oxford, GB), Thomas W. Reps (University of Wisconsin – Madison, US), Sanjit A. Seshia (University of California – Berkeley, US), Aditya Thakur (University of Wisconsin – Madison, US)
August 24–29, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14351

14352 – Next Generation Static Software Analysis Tools
Patrick Cousot (ENS – Paris, FR), Klaus Havelund (CalTech – Pasadena, US), Daniel Kroening (University of Oxford, GB), Carsten Sinz (KIT – Karlsruher Institut für Technologie, DE), Mana Taghdiri (KIT – Karlsruher Institut für Technologie, DE)
August 24–29, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14352

14361 – Computational Aspects of Fabrication
Marc Alexa (TU Berlin, DE), Bernd Bickel (Disney Research – Zürich, CH), Sara McMains (University of California – Berkeley, US), Holly E. Rushmeier (Yale University, US)
August 31 to September 5, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14361

14362 – Augmenting Human Memory – Capture and Recall in the Era of Lifelogging
Mark Billinghurst (University of Canterbury – Christchurch, NZ), Nigel Davies (Lancaster University, GB), Marc Langheinrich (University of Lugano, CH), Albrecht Schmidt (Universität Stuttgart, DE)
August 31 to September 5, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14362

14371 – Adjoint Methods in Computational Science, Engineering, and Finance
Nicolas R. Gauger (TU Kaiserslautern, DE), Michael Giles (University of Oxford, GB), Max D. Gunzburger (Florida State University, US), Uwe Naumann (RWTH Aachen, DE)
September 7–12, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14371

14372 – Analysis of Algorithms Beyond the Worst Case
Maria-Florina Balcan (Carnegie Mellon University – Pittsburgh, US), Bodo Manthey (University of Twente, NL), Heiko Röglin (Universität Bonn, DE), Tim Roughgarden (Stanford University, US)
September 7–12, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14372

14381 – Neural-Symbolic Learning and Reasoning
Artur d'Avila Garcez (City University – London, GB), Marco Gori (University of Siena, IT), Pascal Hitzler (Wright State University – Dayton, US), Luis Lamb (Federal University of Rio Grande do Sul, BR)
September 14–19, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14381

14391 – Algebra in Computational Complexity
Manindra Agrawal (IIT – Kanpur, IN), Valentine Kabanets (Simon Fraser University – Burnaby, CA), Thomas Thierauf (Hochschule Aalen, DE), Christopher Umans (CalTech – Pasadena, US)
September 21–26, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14391

14402 – Resilience in Exascale Computing
Hermann Härtig (TU Dresden, DE), Satoshi Matsuoka (Tokyo Institute of Technology, JP), Frank Mueller (North Carolina State University, US), Alexander Reinefeld (Konrad-Zuse-Zentrum – Berlin, DE)
September 28 to October 1, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14402

14411 – Constraints, Optimization and Data
Luc De Raedt (KU Leuven, BE), Siegfried Nijssen (KU Leuven, BE), Barry O'Sullivan (University College Cork, IE), Michele Sebag (University of Paris South XI, FR), Pascal Van Hentenryck (The University of Melbourne, AU)
October 5–10, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14411

14412 – Globalizing Domain-Specific Languages
Betty H. C. Cheng (Michigan State University – East Lansing, US), Benoit Combemale (IRISA / CNRS, FR), Robert B. France (Colorado State University, US), Jean-Marc Jezequel (IRISA / CNRS, FR), Bernhard Rumpe (RWTH Aachen, DE)
October 5–10, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14412

14421 – Optimal algorithms and proofs
Olaf Beyersdorff (University of Leeds, GB), Edward A. Hirsch (Steklov Institute – St. Petersburg, RU), Jan Krajicek (Charles University – Prague, CZ), Rahul Santhanam (University of Edinburgh, GB)
October 12–17, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14421

14441 – Modeling, Verification, and Control of Complex Systems for Energy Networks
Alessandro Abate (University of Oxford, GB), Martin Fränzle (Universität Oldenburg, DE), Ian Hiskens (University of Michigan – Ann Arbor, US), Martin Strelec (UWB – Pilsen, CZ)
October 26–31, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14441

14442 – Symbolic Execution and Constraint Solving
Cristian Cadar (Imperial College London, GB), Vijay Ganesh (University of Waterloo, CA), Raimondas Sasnauskas (University of Utah, US), Koushik Sen (University of California – Berkeley, US)
October 27–30, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14442

14451 – Optimality and tight results in parameterized complexity
Stefan Kratsch (TU Berlin, DE), Daniel Lokshtanov (University of Bergen, NO), Dániel Marx (Hungarian Academy of Sciences – Budapest, HU), Peter Rossmanith (RWTH Aachen, DE)
November 2–7, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14451

14452 – Algorithmic Cheminformatics
Wolfgang Banzhaf (Memorial University of Newfoundland, CA), Christoph Flamm (Universität Wien, AT), Daniel Merkle (University of Southern Denmark – Odense, DK), Peter F. Stadler (Universität Leipzig – IZBI, DE)
November 2–7, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14452

14461 – High-performance Graph Algorithms and Applications in Computational Science
Ulrich Carsten Meyer (Goethe-Universität Frankfurt am Main, DE), Henning Meyerhenke (KIT – Karlsruher Institut für Technologie, DE), Ali Pinar (Sandia Nat. Labs – Livermore, US), Ilya Safro (Clemson University, US)
November 9–14, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14461

14462 – Systems and Algorithms for Large-scale Graph Analytics
Derek Murray (, US), Amitabha Roy (EPFL – Lausanne, CH), Eiko Yoneki (University of Cambridge, GB)
November 9–12, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14462

14471 – Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment
Jon Crowcroft (University of Cambridge, GB), Arjuna Sathiaseelan (University of Cambridge, GB), Adam Wolisz (TU Berlin, DE)
November 16–21, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14471

14481 – Multiscale Spatial Computational Systems Biology
David Gilbert (Brunel University, GB), Monika Heiner (BTU Cottbus, DE), Koichi Takahashi (Osaka University, JP), Adelinde M. Uhrmacher (Universität Rostock, DE)
November 23–28, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14481

14482 – Automated Planning and Model Checking
Alessandro Cimatti (Bruno Kessler Foundation – Trento, IT), Stefan Edelkamp (Universität Bremen, DE), Maria Fox (King's College London, GB), Daniele Magazzeni (King's College London, GB), Erion Plaku (CUA – Washington, US)
November 23–28, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14482

14491 – Socio-Technical Security Metrics
Dieter Gollmann (TU Hamburg-Harburg, DE), Cormac Herley (Microsoft Corp. – Redmond, US), Vincent Koenig (University of Luxembourg, LU), Wolter Pieters (TU Delft & University of Twente, NL), Martina Angela Sasse (University College London, GB)
November 30 to December 5, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14491

14492 – The synergy between programming languages and cryptography
Gilles Barthe (IMDEA Software – Madrid, ES), Michael Hicks (University of Maryland, US), Florian Kerschbaum (SAP AG – Karlsruhe, DE), Dominique Unruh (University of Tartu, EE)
November 30 to December 5, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14492

14511 – Programming Languages for Big Data (PlanBig)
James Cheney (University of Edinburgh, GB), Torsten Grust (Universität Tübingen, DE), Dimitrios Vytiniotis (Microsoft Research UK – Cambridge, GB)
December 14–19, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14511

14512 – Collective Adaptive Systems: Qualitative and Quantitative Modelling and Analysis
Jane Hillston (University of Edinburgh, GB), Matthias Hölzl (LMU München, DE), Jeremy Pitt (Imperial College London, GB), Martin Wirsing (LMU München, DE), Franco Zambonelli (University of Modena, IT)
December 14–19, 2014 | Dagstuhl Seminar | http://www.dagstuhl.de/14512

Dagstuhl-Perspektiven-Workshops 13.2 *Dagstuhl Perspectives Workshops*

14022 – Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing
Peer-Timo Bremer (LLNL – Livermore, US), Bernd Mohr (Jülich Supercomputing Centre, DE), Valerio Pascucci (University of Utah, US), Martin Schulz (LLNL – Livermore, US)
January 5–10, 2014 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14022

14112 – Massive Open Online Courses: Current State and Perspectives
Pierre Dillenbourg (EPFL – Lausanne, CH), Claude Kirchner (INRIA – Le Chesnay, FR), John C. Mitchell (Stanford University, US), Martin Wirsing (LMU München, DE)
March 10–13, 2014 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14112

14182 – Categorical Methods at the Crossroads
Samson Abramsky (University of Oxford, GB), John C. Baez (University of California – Riverside, US), Fabio Gadducci (University of Pisa, IT), Viktor Winschel (Universität Mannheim, DE)
April 27 to May 2, 2014 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14182

14272 – Exploring Interdisciplinary Grand Challenges in ICT Design to Support Proactive Health and Wellbeing
Lars L. Andersen (NRCWE – Copenhagen, DK), Elizabeth F. Churchill (eBay Research Labs, US), m.c. schraefel (University of Southampton, GB)
June 29 to July 2, 2014 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14272

14401 – Privacy and Security in an Age of Surveillance
Matt Blaze (University of Pennsylvania, US), Bart Preneel (KU Leuven, BE), Phillip Rogaway (University of California – Davis, US), Mark D. Ryan (University of Birmingham, GB), Peter Y. A. Ryan (University of Luxembourg, LU)
September 28 to October 2, 2014 | Dagstuhl Perspectives Workshop | http://www.dagstuhl.de/14401

GI-Dagstuhl-Seminare 13.3 *GI-Dagstuhl Seminars*

14382 – Control Theory meets Software Engineering
Antonio Filieri (Universität Stuttgart, DE), Martina Maggio (Lund University, SE)
September 14–19, 2014 | Meeting | http://www.dagstuhl.de/14382

14433 – Software Engineering for Self-Adaptive Systems
Alessandra Gorla (Universität des Saarlandes, DE), Matthias Tichy (Chalmers UT – Göteborg, SE), Thomas Vogel (Hasso-Plattner-Institut – Potsdam, DE)
October 19–24, 2014 | Meeting | http://www.dagstuhl.de/14433

14502 – Informatik@Schule – Agenda für informatische Bildung in der Schule
Rainer Gemulla (Universität Mannheim, DE), Wim Martens (Universität Bayreuth, DE), Johannes Schöning (Hasselt University – Diepenbeek, BE), Carsten Schulte (FU Berlin, DE)
December 7–10, 2014 | Meeting | http://www.dagstuhl.de/14502

Lehrveranstaltungen 13.4 *Educational Events*

14152 – Spring School: "Individual-centered Approaches to Speech Processing"
Bernd Möbius (Universität des Saarlandes, DE), Jürgen Trouvain (Universität des Saarlandes, DE)
April 6–9, 2014 | Meeting | http://www.dagstuhl.de/14152

14222 – Workshop Wissenschaftsjournalismus
Roswitha Bardohl (Schloss Dagstuhl, DE), Gordon Bolduan (Universität des Saarlandes, DE), Tim Schröder (Oldenburg, DE)
May 25–28, 2014 | Meeting | http://www.dagstuhl.de/14222

14392 – Autumn School 2014 for Information Retrieval and Information Foraging
Ingo Frommholz (University of Bedfordshire – Luton, GB), Norbert Fuhr (Universität Duisburg-Essen, DE), Thomas Mandl (Universität Hildesheim, DE)
September 21–26, 2014 | Meeting | http://www.dagstuhl.de/14392

14503 – Lehrerfortbildung in Informatik
Roswitha Bardohl (Schloss Dagstuhl, DE), Manuel Garcia Mateos (LPM Saarbrücken, DE), Martin Zimnol (Pädagogisches Landesinstitut Rheinland-Pfalz, DE)
December 10–12, 2014 | Meeting | http://www.dagstuhl.de/14503

Sonstige Veranstaltungen 13.5 *Other Events*

14033 – Lehrstuhltreffen AG Schneider / Sturm
Jörn Schneider (Hochschule Trier, DE)
January 15–17, 2014 | Meeting | http://www.dagstuhl.de/14033

14103 – Modellbasierte Entwicklung eingebetteter Systeme (MBEES)
Bernhard Schätz (fortiss GmbH – München, DE)
March 5–7, 2014 | Meeting | http://www.dagstuhl.de/14103

14153 – International Workshop: "Multilinguality in Speech Research: Data, Methods and Models"
Bernd Möbius (Universität des Saarlandes, DE), Jürgen Trouvain (Universität des Saarlandes, DE)
April 9–11, 2014 | Meeting | http://www.dagstuhl.de/14153

14154 – Lehrstuhltreffen AG Zeller
Andreas Zeller (Universität des Saarlandes, DE)
April 6–8, 2014 | Meeting | http://www.dagstuhl.de/14154

14155 – Klausurtagung Telematik Karlsruhe
Sören Finster (KIT – Karlsruher Institut für Technologie, DE), Martina Zitterbart (KIT – Karlsruher Institut für Technologie, DE)
April 9–11, 2014 | Meeting | http://www.dagstuhl.de/14155

14162 – GIBU 2014: GI-Beirat der Universitätsprofessoren
Gregor Snelting (KIT – Karlsruher Institut für Technologie, DE)
April 13–15, 2014 | Meeting | http://www.dagstuhl.de/14162

14163 – Klausurtagung "LST Halang"
Wolfgang A. Halang (FernUniversität in Hagen, DE)
April 13–16, 2014 | Meeting | http://www.dagstuhl.de/14163

14173 – Gruppentreffen
Hans Hagen (TU Kaiserslautern, DE)
April 24–25, 2014 | Meeting | http://www.dagstuhl.de/14173

14192 – Kolloquium zum GI Dissertationspreis 2013
Steffen Hölldobler (TU Dresden, DE)
May 4–7, 2014 | Meeting | http://www.dagstuhl.de/14192

14193 – FOSD-Treffen
Sven Apel (Universität Passau, DE), Christian Kästner (Carnegie Mellon University, US), Sandro Schulze (TU Braunschweig, DE)
May 4–7, 2014 | Meeting | http://www.dagstuhl.de/14193

14212 – Klausurtagung Informatik THM
Timo Péus (THM – Friedberg, DE)
May 21–23, 2014 | Meeting | http://www.dagstuhl.de/14212

14242 – Lehrstuhltreffen "Embedded Intelligence"
Bernhard Sick (Universität Kassel, DE)
June 9–13, 2014 | Meeting | http://www.dagstuhl.de/14242

14243 – Lehrstuhltreffen Rechtsinformatik
Christoph Sorge (Universität des Saarlandes – Saarbrücken, DE)
June 11–13, 2014 | Meeting | http://www.dagstuhl.de/14243

14252 – Gemeinsamer Workshop der Graduiertenkollegs: Interdisciplinary scientific working principles
Björn Engelmann (Universität Oldenburg, DE), Nils-Erik Flick (Universität Oldenburg, DE), Yang Gao (Universität Oldenburg, DE), Sebastian Hahn (Hasso-Plattner-Institut – Potsdam, DE), Anja Jentzsch (Hasso-Plattner-Institut – Potsdam, DE), Heinrich Ody (Universität Oldenburg, DE), Sebastian Pasewaldt (Hasso-Plattner-Institut – Potsdam, DE)
June 15–18, 2014 | Meeting | http://www.dagstuhl.de/14252

14262 – OPTET Workshop
Sachar Paulus (FH Brandenburg an der Havel, DE)
June 22–25, 2014 | Meeting | http://www.dagstuhl.de/14262

14263 – Workshop Buchprojekt "Corporate Semantic Web"
Anatol Reibold (OntoPort Group – Darmstadt, DE)
June 25–27, 2014 | Meeting | http://www.dagstuhl.de/14263

14283 – Lehrstuhltreffen AG Grimm
Rüdiger Grimm (Universität Koblenz-Landau, DE)
July 9–11, 2014 | Meeting | http://www.dagstuhl.de/14283

14309 – Forschungsaufenthalt
Frances A. Rosamond (Charles Darwin University – Darwin, AU)
July 20–26, 2014 | Meeting | http://www.dagstuhl.de/14309

14339 – Forschungsaufenthalt
Frances A. Rosamond (Charles Darwin University – Darwin, AU)
August 10–14, 2014 | Meeting | http://www.dagstuhl.de/14339

14354 – Klausurtagung Universität des Saarlandes
Volker Linneweber (Universität des Saarlandes, DE)
August 27–28, 2014 | Meeting | http://www.dagstuhl.de/14354

14359 – Forschungsaufenthalt
Frances A. Rosamond (Charles Darwin University – Darwin, AU)
August 24–29, 2014 | Meeting | http://www.dagstuhl.de/14359

14403 – Research Colloquium Automation and Energy Systems Group Fall 2014
Georg Frey (Universität des Saarlandes, DE)
October 1–2, 2014 | Meeting | http://www.dagstuhl.de/14403

14422 – DDI: Facilitating Process and Metadata-Driven Automation in the Social, Economic, and Behavioural Sciences with the Data Documentation Initiative
Arofan Gregory (Open Data Foundation – Tucson, US), Jon Johnson (Centre for Longitudinal Studies London, GB), Wendy Thomas (Univ. of Minnesota – Minneapolis, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 12–17, 2014 | Meeting | http://www.dagstuhl.de/14422

14432 – DDI Lifecycle: Moving Forward
Arofan Gregory (Open Data Foundation – Tucson, US), Larry Hoyle (The University of Kansas – Lawrence, US), Wendy Thomas (Univ. of Minnesota – Minneapolis, US), Mary Vardigan (University of Michigan – Ann Arbor, US), Joachim Wackerow (GESIS – Mannheim, DE)
October 19–24, 2014 | Meeting | http://www.dagstuhl.de/14432

14463 – IFIP TC6 2014/2 Strategic Review Meeting
Jordi Domingo-Pascual (UPC – Barcelona, ES), Gunnar Karlsson (KTH Royal Institute of Technology, SE), Siraj Ahmed Shaikh (Coventry University, GB), Otto Spaniol (RWTH Aachen, DE)
November 12–14, 2014 | Meeting | http://www.dagstuhl.de/14463

14472 – Klausurtagung
Nicola Wolpert (University of Applied Sciences – Stuttgart, DE)
November 16–19, 2014 | Meeting | http://www.dagstuhl.de/14472

14473 – Klausurtagung "LST Schmeck"
Florian Allerding (KIT – Karlsruher Institut für Technologie, DE), Hartmut Schmeck (KIT – Karlsruher Institut für Technologie, DE)
November 19–21, 2014 | Meeting | http://www.dagstuhl.de/14473

14474 – Secan Lab Seminar
Thomas Engel (University of Luxembourg, LU)
November 17–18, 2014 | Meeting | http://www.dagstuhl.de/14474

14475 – SFB-Treffen
Frank Mücklich (Universität des Saarlandes, DE)
November 20–21, 2014 | Meeting | http://www.dagstuhl.de/14475

14504 – Young Researcher's Training SFB-TRR55 Lattice QCD
Sara Collins (Universität Regensburg, DE), Andreas Frommer (Universität Wuppertal, DE), Francesco Knechtli (Universität Wuppertal, DE)
December 7–9, 2014 | Meeting | http://www.dagstuhl.de/14504

Mitglied der Leibniz-Gemeinschaft

Oktavie-Allee, 66687 Wadern, Deutschland

Jahresbericht / Annual Report 2014 | ISSN 2199-1995
http://www.dagstuhl.de