DAGSTUHL
REPORTS

Volume 6, Issue 7, July 2016

Dagstuhl Reports, Vol. 6, Issue 7 ISSN 2192-5283

ISSN 2192-5283

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at
http://www.dagstuhl.de/dagpub/2192-5283

Publication date
November, 2016

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at `http://dnb.d-nb.de`.

Digital Object Identifier: 10.4230/DagRep.6.7.i

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
http://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 16271

Algorithmic Foundations of Programmable Matter

Edited by

Sándor Fekete[1], Andréa W. Richa[2], Kay Römer[3], and Christian Scheideler[4]

1 **TU Braunschweig, DE,** s.fekete@tu-bs.de
2 **Arizona State University – Tempe, US,** aricha@asu.edu
3 **TU Graz, AT,** roemer@tugraz.at
4 **Universität Paderborn, DE,** scheideler@upb.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 16271 "Algorithmic Foundations of Programmable Matter", a new and emerging field that combines theoretical work on algorithms with a wide spectrum of practical applications that reach all the way from small-scale embedded systems to cyber-physical structures at nano-scale.

The aim of the Dagstuhl seminar was to bring together researchers from the algorithms community with selected experts from robotics and distributed systems in order to set a solid base for the development of models, technical solutions, and algorithms that can control programmable matter. Both communities benefited from such a meeting for the following reasons:

- Meeting experts from other fields provided additional insights, challenges and focus when considering work on programmable matter.
- Interacting with colleagues in a close and social manner gave many starting points for continuing collaboration.
- Getting together in a strong, large and enthusiastic group provided the opportunity to plan a number of followup activities.

In the following, we provide details and activities of this successful week.

Seminar July 3–8, 2016 – http://www.dagstuhl.de/16271

1998 ACM Subject Classification C.2.4 [Computer Communications Networks] Distributed Systems, F.1.2 [Computation by Abstract Devices] Modes of Computation, F.2.2 [Analysis of Algorithms and Problem Complexity] Nonnumerical Algorithms and Problems, I.2.9 [Artificial Intelligence] Robotics, I.2.9 [Artificial Intelligence] Robotics – Modular Reconfigurable Robots, I.6 [Simulation and Modeling] General

Keywords and phrases distributed algorithms, distributed systems, programmable matter, robotics, self-organization

Digital Object Identifier 10.4230/DagRep.6.7.1

Edited in cooperation with Thim Strothmann

Algorithmic Foundations of Programmable Matter, *Dagstuhl Reports*, Vol. 6, Issue 7, pp. 1–14
Editors: Sándor Fekete, Andréa W. Richa, Kay Römer, and Christian Scheideler
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

Sándor Fekete
Andréa W. Richa
Kay Römer
Christian Scheideler

Programmable matter refers to a substance that has the ability to change its physical properties (shape, density, moduli, conductivity, optical properties, etc.) in a programmable fashion, based upon user input or autonomous sensing. The potential applications are endless, e.g., smart materials, autonomous monitoring and repair, or minimal invasive surgery. Thus, there is a high relevance of this topic to industry and society in general, and much research has been invested in the past decade to fabricate programmable matter. However, fabrication is only part of the story: without a proper understanding of how to program that matter, complex tasks such as minimal invasive surgery will be out of reach. Unfortunately, only very few people in the algorithms community have worked on programmable matter so far, so programmable matter has not received the attention it deserves given the importance of that topic.

The Dagstuhl seminar "Algorithmic Foundations of Programmable Matter" aimed at resolving that problem by getting together a critical mass of people from algorithms with a selection of experts from distributed systems and robotics in order to discuss and develop models, algorithms, and technical solutions for programmable matter.

The aim of the proposed seminar was to bring together researchers from the algorithms community with selected experts from robotics and distributed systems in order to set a solid base for the development of models, technical solutions, and algorithms that can control programmable matter. The overall mix worked quite well: researchers from the more practical side (such as Julien Bourgeois, Nikolaus Correll, Ted Pavlic, Kay Römer, among others) interacted well with participants from the theoretical side (e.g., Jennifer Welch, Andrea Richa, Christian Scheideler, Sándor Fekete, and many others). Particularly interesting to see were well-developed but still expanding areas, such as tile self-assembly that already combines theory and practice (with visible and well-connected scientists such as Damien Woods, Matt Patitz, David Doty, Andrew Winslow, Robert Schweller) or multi-robot systems (Julien Bourgeois, Nikolaus Correll, Matteo Lasagni, André Naz, Benoît Piranda, Kay Römer).

The seminar program started with a set of four tutorial talks given by representatives from the different sets of participants to establish a common ground for discussion. From the robotics and distributed system side, Nikolaus Correll and Julien Bourgeois gave tutorials on smart programmable materials and on the claytronics programmable matter framework respectively. From the bioengineering side, Ted Pavlic gave a tutorial on natural systems that may inspire programmable matter. From the algorithmic side, Jacob Hendricks gave a tutorial on algorithmic self-assembly. In the mornings of the remaining four days, selected participants offered shorter presentations with a special focus on experience from the past work and especially also open problems and challenges. Two of the afternoons were devoted to discussions in breakout groups. Four breakout groups were formed, each with less than 10 participants to allow for intense interaction. Inspired by a classification of research questions in biology into "why?" and "how?" questions presented in Ted Pavlic's tutorial, the first breakout session was devoted to the "why?" questions underpinning programmable matter, especially also appropriate models of programmable matter systems (both biological or engineered) suitable for algorithmic research. The second breakout sessions towards the end of the seminar was devoted to a set of specific questions given by the organizers that resulted from the discussions among the participants, they included both research questions and organizational questions (e.g., how to proceed after the Dagstuhl seminar). After each of the two breakout sessions, one participant of each of the four breakout groups reported back the main findings of the discussions to the plenum, leading to further discussion among all participants. One of the afternoons was devoted to a hike to a nearby village, where the participants also visited a small museum devoted to programmable mechanical musical devices.

The seminar was an overwhelming success. In particular, bringing together participants from a number of different but partially overlapping areas, in order to exchange problems and challenges on a newly developing field turned out to be excellent for the setting of Dagstuhl – and the opportunities provided at Dagstuhl are perfect for starting a new community.

Participants were enthusiastic on a number of different levels:

- Meeting experts from other fields provided additional insights, challenges and focus when considering work on programmable matter.
- Interacting with colleagues in a close and social manner gave many starting points for continuing collaboration.
- Getting together in a strong, large and enthusiastic group provided the opportunity to plan a number of followup activities.

The latter include connecting participants via a mailing list, the planning and writing of survey articles in highly visible publication outlets, and a starting point for specific scientific workshops and conferences.

Participants were highly enthusiastic about the possibility of another Dagstuhl workshop in the future; organizers will keep the ball rolling on this – most likely, for an application in the coming spring, so that some more details can be worked out in the meantime.

2 Table of Contents

3 Overview of Talks

3.1 Claytronics: an Instance of Programmable Matter

Julien Bourgeois (FEMTO-ST Institute – Montbéliard, FR)

Joint work of Julien Bourgeois, Seth Copen Goldstein
Main reference J. Bourgeois, S. Copen Goldstein, "Distributed Intelligent MEMS: Progresses and Perspectives", IEEE Systems Journal, 9(3):1057–106, 2015.
URL http://dx.doi.org/10.1109/JSYST.2013.2281124

Programmable matter (PM) has different meanings but they can be sorted depending on four properties: Evolutivity, Programmability, Autonomy and Interactivity. In my talk, I will present the Claytronics project which is an instance of PM, evolutive, programmable, autonmous and interactive. In Claytronics, PM is defined as a huge modular self-reconfigurable robot. To manage the complexity of this kind of environment, we propose a complete environment including programmable hardware, a programming langage, a compiler, a simulator, a debugger and distributed algorithms.

3.2 A Markov Chain Algorithm for Compression in Self-Organizing Particle Systems

Sarah Cannon (Georgia Institute of Technology – Atlanta, US)

Joint work of Sarah Cannon, Joshua J. Daymude, Dana Randall, Andréa W. Richa
Main reference S. Cannon, J. J. Daymude, D. Randall, A. W. Richa, "A Markov Chain Algorithm for Compression in Self-Organizing Particle Systems", in Proc. of the 2016 ACM Symp. on Principles of Distributed Computing (PODC'16), pp. 279–288, ACM, 2016.
URL http://dx.doi.org/10.1145/2933057.2933107

One can model programmable matter as a collection of simple computational elements (called particles) with limited (constant-size) memory that self-organize to solve system-wide problems of movement, configuration, and coordination. In recent work with Joshua J. Daymude, Andrea Richa, and Dana Randall, we focused on the compression problem, in which the particle system gathers as tightly together as possible, as in a sphere or its equivalent in the presence of some underlying geometry. More specifically, we presented a fully distributed, local, and asynchronous algorithm that leads the system to converge to a configuration with small perimeter. Our Markov chain based algorithm solves the compression problem under the geometric amoebot model, for particle systems that begin in a connected configuration with no holes. I will give a brief overview of Markov chains, describe our Markov chain and why it achieves particle compression, and show how it leads to a fully distributed, local, and asynchronous protocol each particle can run independently. Furthermore, I'll discuss how Markov chains might be amenable for use in other programmable matter contexts

3.3 Algorithm design for swarm robotics and smart materials

Nikolaus Correll (University of Colorado – Boulder, US)

Main reference M. A. McEvoy, N. Correll, "Materials that couple sensing, actuation, computation, and communication", Science, 347(6228), p. 1261689, 2015.
URL http://dx.doi.org/10.1126/science.1261689

"Programmable Matter" is a conjunction of a discrete program and continuous matter. Where the line between the two needs to be drawn is currently unclear. One approach is to abstract matter to the point where it can be treated exclusively by discrete models. Another approach is to think about individual elements becoming so small that they are captured by continuous physics such as fluid dynamics. In this tutorial, I argue for and explain a hybrid automata model that consists of a discrete network of computers, which can sense and actuate on a continuous material that the network is integrated in. Here, communication not only happens through the network, but also in the material itself via sensor/actuator coupling. I illustrate this approach using a series of "robotic materials" including a sensing skin, a shape-changing beam, and a modular wall that can recognize gestures. These systems demonstrate a number of algorithmic challenges ranging from networking, distributed control and optimization, and programming. At the same time, each instance illustrates that material properties strongly influence algorithmic design and vice versa.

3.4 Dynamic Networks of Computationally Challenged Devices: the Passive Case

Yuval Emek (Technion – Haifa, IL)

Joint work of Yuval Emek, Jara Uitto, Roger, Wattenhofer
Main reference Y. Emek, R. Wattenhofer, "Stone age distributed computing", in Proc. of the 2013 ACM Symp. on Principles of Distributed Computing (PODC'13), pp. 137–146, ACM, 2013.
URL http://dx.doi.org/10.1145/2484239.2484244

Motivated by applications in biology and nano-technology, the trend of applying the "distributed computing approach" to networks of message passing devices with weak computation (as well as communication) capabilities is gaining momentum. So far, most of the advances have been made under the assumption that (1) the network is static; or (2) the dynamic behavior of the network is dictated by devices that can actively control their own motion. In this talk, I'd like to discuss some research questions related to such networks that undergo dynamic (adversarial) topology changes in which the devices only play a passive role.

3.5 Algorithms for robot navigation: From optimizing individual robots to particle swarms

Sándor Fekete (TU Braunschweig, DE)

Joint work of Aaron Becker, Erik D. Demaine, Maximilian Ernestus, Sándor Fekete, Michael Hemmer, Alexander Kröller, Dominik Krupke, SeoungKyou Lee, James McLurkin, Rose Morris-Wright, Christiane Schmidt, S. H. Mohtasham

Planning and optimizing the motion of one or several robots poses a wide range of problems. How can we coordinate a group of weak robots to explore an unknown environment? How can we ensure that a swarm of very simple robots with local capabilities can deal with conflicting global requirements? And how can a particle swarm perform complex operations? We will demonstrate how an appropriate spectrum of algorithmic methods in combination with geometry can be used to achieve progress on all of these challenges.

3.6 The Amoebot Model

Robert Gmyr (Universität Paderborn, DE)

Joint work of Zahra Derakhshandeh, Robert Gmyr, Andréa W. Richa, Christian Scheideler, Thim Strothmann
Main reference Z. Derakhshandeh, R. Gmyr, A. W. Richa, C. Scheideler, T. Strothmann, "Universal Shape Formation for Programmable Matter", in Proc. of the 28th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'16), pp. 289–299, ACM, 2016.
URL http://dx.doi.org/10.1145/2935764.2935784

We envision programmable matter consisting of systems of computationally limited devices that are able to self-organize in order to achieve a desired collective goal without the need for central control or external intervention. Our formal investigation of programmable matter is based on the Amoebot model. In this talk, I will give a brief introduction to this model. Furthermore, I will give an overview of our work on three central problems, namely shape formation, coating, and leader election.

3.7 Dances with Plants: Robot-supported Programmable Living Matter

Heiko Hamann (Universität Paderborn, DE)

Joint work of Mostafa Wahby, Mohammad Divband Soorati, Heiko Hamann
Main reference H. Hamann, M. Wahby, T. Schmickl, P. Zahadat, D. Hofstadler, K. Støy, S. Risi, A. Faina, F. Veenstra, S. Kernbach, I. Kuksin, O. Kernbach, P. Ayres, P. Wojtaszek, "*flora robotica* – Mixed Societies of Symbiotic Robot-Plant Bio-Hybrids", in Proc. of the IEEE Symp. on Computational Intelligence (SSCI'15), pp. 1102–1109, IEEE, 2015.
URL http://dx.doi.org/10.1109/SSCI.2015.158

Besides standard self-reconfiguring modular robotics and self-assembly, robots can also be mixed with other components to form heterogeneous systems. For example, combining natural plants and distributed robot systems offers new approaches to programmable matter. Instead of applying methods of synthetic biology, the idea here is to make use of natural plant behaviors to control them. A second example is self-organized swarm construction of possibly actuated structures. These approaches offer unique advantages, such as growth of additional material for free, environmental safety, and simplicity, despite their limitations in flexibility concerning possible structures and potential for reconfigurations.

3.8 Introduction to Modeling Algorithmic Self-Assembling Systems

Jacob Hendricks (University of Wisconsin – River Falls, US)

This talk introduces theoretical tile based models of self-assembly. We first give the definition of Winfree's abstract Tile Assembly Model (aTAM) that was developed to study DNA based molecular building blocks. This talk proceeds with examples of specific tile assembly systems such as binary counters and Turing machine simulators as these demonstrate the possibility of algorithmic self-assembly. Then we discuss a few specific topics in the field. These topics include non-cooperative self-assembly, various models of self-assembly, common benchmarks for determining the capabilities and limitations of models, simulation as a means of comparing models, and finally, the notion of intrinsic universality. Topics have been selected to provide a bird's-eye view of a theoretician's considerations about modelling self-assembling systems using tile assembly models.

3.9 Advantages, Limitations, Challenges of Tendon-Driven Programmable Chains

Matteo Lasagni (TU Graz, AT)

Joint work of Matteo Lasagni, Kay Römer
Main reference M. Lasagni, K. Römer, "Force-guiding particle chains for shape-shifting displays", in Proc. of the 2014 IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems (IROS'14), pp. 3912–3918, IEEE, 2014.
URL http://dx.doi.org/10.1109/IROS.2014.6943112

One of the first and most relevant questions when designing a shape-shifting material concerns particles' topology and hence their mobility. We can distinguish between detachable and non-detachable topologies. Detachable topologies allow particles to temporarily detach from the ensemble and freely migrate to different locations in order to obtain a shape-shift. In contrast, non-detachable topologies constrain particles to occupy a fixed location where only relative displacement between adjacent particles and/or particle deformation are allowed for shape-shift. Despite the fact that detachable topologies allow the formation of literally any shape, the complex architecture to enable particle migration, generally consisting of built-in actuators and latching mechanisms, raises costs and limits the scalability of the whole system, and causes inherent problems concerning particle power supply and communication. In our work, we demonstrate how a nondetachable topology can allow the formation of arbitrary complex shapes, thereby avoiding or at least limiting the above-mentioned problems. In particular, scalability and cost-effectiveness derive from the concatenation of semi-active particles without bulky built-in actuators and latches. Such particles, forming piecewise bendable chains, exploit remotely generated forces to self-actuate and hence to control the local curvature of the chain. Multiple chains can be combined to form a shape-shifting surface to support novel applications like 3D tangible displays or programmable molds. One major challenge concerns the actuation of the system. Without the support of optimal planning strategies able to schedule proper particle actuation, unbearable actuation forces might occur, for example, due to inconvenient leverage effects, with negative consequences for the system stability and integrity. Starting from the current configuration and aiming at the final target configuration, optimal planning techniques should explore the large set of possible *next configurations* where only a limited number of particles can actuate, and determine in which cases the intensity of the actuation forces is acceptable. The optimization problem involves not only a single independent chain, but applies simultaneously to all chains. Due to mechanical constraints, indeed, all chains need to actuate the same number of particles at each reconfiguration step. This calls for models able to predict the behavior of the whole system upon the application of specific control input in order to support optimal planning algorithms. Such models need to be sufficiently accurate to be consistent with reality but also computationally efficient to allow planning in reasonable time. An important question concerns the determination of an acceptable trade-off between these two aspects.

3.10 Programmable Matter for Dynamic Environments

Othon Michail (CTI – Rion, GR)

Joint work of Othon Michail, Paul G. Spirakis
Main reference O. Michail, P. G. Spirakis, "Terminating population protocols via some minimal global knowledge assumptions", J. Parallel Distrib. Comput., Vol. 81–82, pp. 1–10, 2015.
URL http://dx.doi.org/10.1016/j.jpdc.2015.02.005

We discuss two recent theoretical models of programmable matter operating in a dynamic environment. In the first model, all devices are finite automata, begin from the same initial state, execute the same protocol, and can only interact in pairs. The interactions are scheduled by a fair (or uniform random) scheduler, in the spirit of Population Protocols. When two devices interact, the protocol takes as input their states and the state of the connection between them (on/off) and updates all of them. Initially all connections are off. The goal of such protocols is to eventually construct a desired stable network, induced by the edges that are on. We present protocols and lower bounds for several basic network construction problems and also universality results. We next discuss a more applied version of this minimal and abstract model, enriched with geometric constraints, aiming at capturing some first physical restrictions in potential future programmable matter systems operating in dynamic environments.

3.11 Energy Harvesting in-vivo Nano-Robots in Caterpillar Swarm

Venkateswarlu Muni (Ben Gurion University of the Negev – Beer Sheva, IL)

Joint work of Shlomi Dolev, Sergey Frenkel, Venkateswarlu Muni, Michael Rosenblit, Ram Prasadh Narayanan

Biological collaborative systems behavior is fascinating, urging researchers to mimic their behavior through programmable matters. These matters constitute a particle system, wherein the particles bind with the neighboring particles to swarm and navigate. Caterpillar swarm inspired particle systems involves layered architecture with single to a predefined number of layers. Through this work, a coordinated layered particle system inspired by caterpillar swarm is discussed. We first propose a novel design for produce-able nano-particles that uses electrodes to harvest electricity from the blood serum, energy that can be later used for swarm inter and/or outer communication, moving, coordination, sensing and acting according to a given (instructing) program. The benefit of moving and acting in a swarm is demonstrated by a design of telescopic movement in pipes (e.g., blood vessels), wherein each layer uses the accumulated speed of all layers below and moves faster, thus, mimicking the faster motion of the caterpillar swarm.

3.12 Algorithmic design of complex 3D DNA origami structures

Pekka Orponen (Aalto University, FI)

Joint work of Erik Benson, Abdulmelik Mohammed, Johan Gardell, Sergej Masich, Eugen Czeizler, Pekka Orponen, Björn Högberg
Main reference E. Benson, A. Mohammed, J. Gardell, S. Masich, E. Czeizler, P. Orponen, B. Högberg, "DNA rendering of polyhedral meshes at the nanoscale", Nature, 523(7561):441–444, 2015.
URL http://dx.doi.org/10.1038/nature14586

In a recent work (Nature 523:441–444, July 2015), we described a general methodology and software pipeline for rendering 3D polyhedral mesh designs in DNA. In this talk, I will first summarise the basic idea of Paul Rothemund's DNA origami technique which also underlies our approach, and then proceed to discuss the graph-theoretic concepts and algorithmic ideas used in extending his technique from 2D patterns to 3D wireframe mesh structures. The reliability and generality of the approach is demonstrated by a number of electron microscopy images of synthesised nanostructures, including a 50-nm rendering of the widely-used Stanford Bunny model. I will also touch on the challenges of using DNA as a substrate for complex designs, and some related open questions.

3.13 Algorithmic Foundations of Biological Matter: Faster, Cheaper, and More Out of Control

Theodore P. Pavlic (Arizona State University – Tempe, US)

For at least the past 30 years, there has been much interest in developing programmable matter solutions that have been inspired by related phenomena in nature. Visionaries in computer science from the 1980's promised that such life-like phenomena would be possible, and yet the programmable matter of today still has limited capabilities. Rather than groups of relatively inexpensive agents grouping together to form an intelligent and flexible collective matter, the automation systems we have today combine relatively intelligent individual units together into rigid and often static structures that have no ability to adapt to the surrounding environment.

To make real progress in understanding the algorithms responsible for nature's success, computer scientists and engineers need to become familiar with the taxonomy of scientific questions in biology. The outputs of biological evolution are shaped not only by adaptive value (i.e., a design objective), but also by phylogeny (i.e., structures inherited from earlier forms), ontogeny (i.e., the process of constructing the object), and the actual mechanism of action that interacts with the environment. The former of these two pressures – adaptive value and phylogeny – are the subject of the "Why" questions of biology, and those questions must always be conditioned by the ancestral environment. The latter two of these pressures – ontogeny and mechanism – are the subject of the "How" questions of biology, and those questions must always be conditioned by the modern environment. This characterization of biological questions is not unlike the ways in which computer scientists consider algorithms and their implementations – in terms of design objectives, platforms, and algorithms that operate on those platforms. However, taking the objective–platform–algorithm approach alone with biological systems obscures details about evolution and ecology that are necessary for understanding how a biological system could possibly be working and whether it is really appropriate to take such an approach with an engineered system. Additionally, taking the effort to understand adaptive value of biological phenomena can provide interesting new motivations for problems that could be solved in engineered systems, albeit using totally different mechanisms.

When learning about programmable matter from biological systems, there are some biological model systems that are better matches than others. A flock of starlings produces beautiful patterns of so-called "active matter" in three-dimensional space. However, such patterns likely have no adaptive value themselves and are simply an epiphenomenon of behaviors which were shaped by individual-level selection to reduce predation risk. As a consequence, anyone trying to reverse engineer such patterns will likely not find much insight in biology. However, army ant bivouacs that are large, self-healing, self-regulating balls of ants in the middle of tropical forest may be a good model as this group-level phenomena is likely under selective pressure by natural selection. Similarly, army ant bridges that move and grow to some intermediate length have been shown to likely provide adaptive benefit to the group. So army ants seem to be a more fitting model for programmable matter than starling flocks. Ants, in general, are attractive targets as all ants evolved from an ancestor that was a solitary wasp; in essence, ants are distributed wasps, with individual workers taking on specialized tasks and aggregating together to form a colony that does what was once the work of a single individual. Although this description of ants is relatively general, there are a wide variety of different problem-solving methods that have evolved across ant taxa. Ants that use pheromone trails as external spatial shared memory apparently vary in how important those trails factor in their individual behaviors, and those variances allow some species of ants to be able to better track changes in the environment than other ants (at the cost of not being able to converge on group-level decisions as quickly). Ants that do not use pheromone trails make foraging and nest-site selection decisions much more slowly, but they have a notable ability to aggregate information from apparently irrational individuals that still produces rational outcomes at the group level. Laboratory tests that induce irrationality in individuals fail to generate the same irrationality in the groups. This kind of aggregation process would be valuable in the development of programmable matter that gathers information from its environment and processes it in a decentralized manner.

Beyond ants, there are other interesting systems that may be good models for programmable matter. One example is the slime mould, in particular the multinucleate slime mould. The multinucleate slime mould is a single cell, formed by the fusion of multiple cells, filled with nuclei and other organelles and surrounded by a flexible, growing membrane that gives the whole unit amoeboid characteristics. Through a decentralized process that dynamically induces pressure gradients around the amoeba-like macroscopic cell, it grows and shrinks and can perform a variety of interesting computations – like finding the shortest path through a maze, determining the right path through a Towers of Hanoi decision tree, and even re-allocating its biomass across different food sources with different mixtures of macronutrients so that the combined intake of those macronutrients is regulated to set levels. While both ants and slime moulds appear somewhat social, it is rare to think about the process of development as a decentralized, collective behavior or computation. However, developmental biologists have shown that cells are effectively executing programs based on internal state and local information from their microenvironment. Following these ideas, some have found ways to re-grow limbs on vertebrates and even induce multiple heads to grow on invertebrates. In the case of the example of the already regenerative planaria (flat worm), modifications induced by a temporary stimulus are latched into the tissue and remembered in all future generations. Those modifications can only be erased by applying drugs used in humans for suppressing memories. This suggests that the somatic tissue all around the body is itself a primitive neural network, capable of information processing and storage. Thus, biological tissue may be a very fitting model system for programmable matter.

In this talk, I will elaborate on these ideas and provide more examples of other biological systems that may provide useful insights when designing programmable matter that may someday finally realize a dream that has been deferred for decades.

3.14 VisibleSim:Your simulator for Programmable Matter

Benoît Piranda (FEMTO-ST Institute – Montbéliard, FR)

Joint work of Julien Bourgeois, Dominique Dhoutaut, André Naz, Benoît Piranda, Pierre Thalamy, Thadeu Tucci
Main reference D. Dhoutaut, B. Piranda, J. Bourgeois, "Efficient simulation of distributed sensing and control environments", in Proc. of the 2013 IEEE Int'l Conf. on Internet of Things (iThings'13), pp. 452–459, IEEE, 2013.
URL http://dx.doi.org/10.1109/GreenCom-iThings-CPSCom.2013.93

VisibleSim is a 3D simulator for distributed robots in a simulated environments. I propose a short tutorial to write a first distributed code for VisibleSim.

3.15 On obliviousness

Nicola Santoro (Carleton University – Ottawa, CA)

Main reference P. Flocchini, G. Prencipe, N. Santoro, "Distributed Computing by Oblivious Mobile Robots", Synthesis Lectures on Distributed Computing Theory, Vol. 3, No. 2 , pp. 1-185, 2012.
URL http://dx.doi.org/10.2200/S00440ED1V01Y201208DCT010

The presence of some form of persistent memory (albeit small in size) is typically assumed in "micro-level" computations (e.g., programmable matter). In contrast, obliviousness (i.e., total absence of persistent memory) is a common restrictions in "macro-level" computations (e.g, autonomous mobile robots) in which I have been involved. On this regards, there are two interesting research questions I would like to share:

- Are meaningful oblivious computations possible at the micro-level?
- What are the precise limits of near-obliviousness (i.e., memory-size thresholds)?

3.16 Theory and practice of large scale molecular-robotic reconfiguration

Damien Woods (California Institute of Technology – Pasadena, US)

Joint work of Ho-Lin Chen, Moya Chen, David Doty, Scott Goodfriend, Nadine Dabby, Dhiraj Holden, Chris Thachuk, Erik Winfree, Damien Woods, Doris Xin, Chun-Tao Yang, Peng Yin

The talk discussed the theory and practice of molecular robotics. I am interested in large scale molecular reconfiguration, where dynamic self-assembling nanostructures change their shape in response to environmental stimuli. I'm very much interested in models that have the potential to be implemented in DNA: a shockingly-well understood and predictable material for nanoscale self-assembly. Specifically, the talk focused attention on questions on the Nubot model and on our initial progress on implementing this style of molecular robotics in the wet-lab.

3.17 Distributed coordination of mobile robots in 3D-space

Yukiko Yamauchi (Kyushu University – Fukuoka, JP)

Joint work of Taichi Uehara, Masafumi Yamashita
Main reference Y. Yamauchi, T. Uehara, M. Yamashita, "Brief Announcement: Pattern Formation Problem for Synchronous Mobile Robots in the Three Dimensional Euclidean Space", in Proc. of the 2016 ACM Symp. on Principles of Distributed Computing (PODC'16), pp. 447–449, ACM, 2016.
URL http://dx.doi.org/10.1145/2933057.2933063

We consider a swarm of autonomous mobile robots moving in the three-dimensional space (3D-space). Each robot is anonymous, oblivious (memory-less), and has neither any access to the global coordinate system nor any communication medium. Many researchers have considered formation problems (point formation, circle formation, pattern formation, etc.) in 2D-space, and it has been shown that the symmetry among the robots determines the patterns that the robots can form. We would like to present our recent results on formation problems in 3D-space, where we encounter rich symmetry represented by rotation groups.

Participants

- Luca Becchetti
Sapienza University of Rome, IT
- Julien Bourgeois
FEMTO-ST Institute – Montbéliard, FR
- Sarah Cannon
Georgia Institute of Technology – Atlanta, US
- Ioannis Chatzigiannakis
Sapienza University of Rome, IT
- Nikolaus Correll
Univ. of Colorado – Boulder, US
- David Doty
Univ. of California – Davis, US
- Yuval Emek
Technion – Haifa, IL
- Sándor Fekete
TU Braunschweig, DE
- Robert Gmyr
Universität Paderborn, DE
- Heiko Hamann
Universität Paderborn, DE
- Jacob Hendricks
University of Wisconsin – River Falls, US
- Stephan Holzer
MIT – Cambridge, US
- Irina Kostitsyna
Free University of Brussels, BE
- Dominik Krupke
TU Braunschweig, DE
- Fabian Daniel Kuhn
Universität Freiburg, DE
- Matteo Lasagni
TU Graz, AT
- Othon Michail
CTI – Rion, GR
- Venkateswarlu Muni
Ben Gurion University of the Negev – Beer Sheva, IL
- André Naz
FEMTO-ST Institute – Montbéliard, FR
- Pekka Orponen
Aalto University, FI
- Matthew J. Patitz
University of Arkansas – Fayetteville, US
- Theodore P. Pavlic
Arizona State University – Tempe, US
- Benoît Piranda
FEMTO-ST Institute – Montbéliard, FR
- Andréa W. Richa
Arizona State University – Tempe, US
- Kay Römer
TU Graz, AT
- Trent Rogers
University of Arkansas – Fayetteville, US
- Nicola Santoro
Carleton Univ. – Ottawa, CA
- Christian Scheideler
Universität Paderborn, DE
- Arne Schmidt
TU Braunschweig, DE
- Robert Schweller
The University of Texas – Pan American – Edinburg, US
- Thim Frederik Strothmann
Universität Paderborn, DE
- Sebastian von Mammen
Universität Augsburg, DE
- Jennifer L. Welch
Texas A&M University – College Station, US
- Andrew Winslow
Free University of Brussels, BE
- Damien Woods
California Institute of Technology – Pasadena, US
- Yukiko Yamauchi
Kyushu Univ. – Fukuoka, JP

Report from Dagstuhl Seminar 16281

Network Latency Control in Data Centres

Edited by

Mohammad Alizadeh Attar[1], Jon Crowcroft[2], Lars Eggert[3], and Klaus Wehrle[4]

1 **MIT – Cambridge, US,** alizadeh@csail.mit.edu
2 **University of Cambridge, GB,** jon.crowcroft@cl.cam.ac.uk
3 **NetApp Deutschland GmbH – Kirchheim, DE,** lars@netapp.com
4 **RWTH Aachen, DE,** wehrle@comsys.rwth-aachen.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 16281 "Network Latency Control in Data Centres". This seminar explored existing and future techniques for controlling data centre latency and thus explores research directions in the new field of data centre latency control in networking research. This need for a new research direction is motivated by the fact that traditional networking equipment and TCP-IP stacks were designed for wide-area networks, where the goal is to maximize throughput, and the control loop between end systems is measured in 10 s of milliseconds. Consequently, this seminar discussed new research direction for data center latency control across the entire software and hardware stack, including in-network solutions, end-host solutions, and others.

Seminar July 10–13, 2016 – http://www.dagstuhl.de/16281
1998 ACM Subject Classification C.2.1 Network Architecture and Design, C.2.2 Network Protocols, C.2.3 Network Operations, D.4.4 Communications Management
Keywords and phrases data centres, end-to-end transport protocols., latency, network architecture, resource control, scheduling
Digital Object Identifier 10.4230/DagRep.6.7.15
Edited in cooperation with Oliver Hohlfeld

1 Executive Summary

Oliver Hohlfeld
Mohammad Alizadeh Attar
Jon Crowcroft
Lars Eggert
Klaus Wehrle

Data centres are at the heart of the modern Internet. They host web services, social networking, cloud computing and are increasingly used by operators to host virtual network functions. All these services have one thing in common: they require extremely low latency communication in the data centre. Consequently we have seen the birth of a new field in networking research – data centre latency control.

Unlike the earlier generation of high-performance computing clusters, data centres have tended to use commodity off-the-shelf servers and switches, and run standard operating systems. However, traditional networking equipment and TCP-IP stacks were designed for wide-area networks, where the goal is to maximize throughput, and the control loop between end systems is measured in 10 s of milliseconds. By contrast, data centres operate on timescales that are several orders of magnitude lower. And while throughput is important, the plentiful bandwidth of data centre networks makes throughput a secondary concern to latency.

This seminar explored existing and future techniques for controlling data centre latency across the entire software and hardware stack, including in-network solutions, end-host solutions, and others. The aims of the seminar are to foster closer collaboration between academic researchers, industry, and operators. 38 researchers attended the multidisciplinary seminar. Over the course of the 3-day seminar, seven presentations were given on various aspects of data center networking. Taking the presentations as input, the workshop then broke into six working groups to discuss research aspects of latency control. The seminar was concluded by voting and discussing on possible conclusions from our discussions. Each conclusion was discussed briefly, then voted on. The outcome of the breakout session as well as the concluding statements are summarized in this report.

2 Table of Contents

3 Overview of Talks

3.1 Data Centre to the Home; Low Delay for All

Bob Briscoe (Simula Research Laboratory – Lysaker, NO)

This talk is about extending the applicability of technology for reducing queuing delay from private data centres to heterogenous data centres and to the wider Internet. Data Centre TCP (DCTCP) can be thought of as an example of how we would now design TCP congestion control given our improvements in understanding since TCP was originally implemented in 1988. As flow rates have increased, the variance of the well-known sawtoothing rate of TCP has grown proportionately, causing large fluctuations in queuing delay. DCTCP solves this problem, so it can consistently keep queuing delay extremely low, whatever the flow rate. DCTCP is therefore categorised as a 'Scalable' congestion control. However, DCTCP requires changes to all senders, all receivers and all switches in a network, therefore (until now) it has not been feasible to deploy it on the public Internet, particularly because it is very aggressive, so existing 'Classic' TCP algorithms tend to starve themselves when they occupy the same queue.

This talk is about a simple technology called the Coupled Dual Queue that solves this coexistence problem without inspecting flows or anything deeper than the outer headers of packets or frames. It allows Scalable congestion controls to coexist with existing Classic controls. It is like a semi-permeable membrane that isolates the new scalable flows from the queuing delay of classic traffic, but the two share the capacity of a link roughly equally as if they were both the same type of TCP sharing a single queue. With this advance, we show that all traffic on the public Internet could enjoy consistently low queuing delay. This should enable new applications that require natural interaction over public networks, such as offloading interactive video processing to cloud data centres. The technology also enables the low queuing delay of scalable congestion controls like DCTCP to be deployed in multi-tenant data centres, or networks of data centres, where there are too many independent administrators to arrange a simultaneous 'flag-day' deployment.

3.2 When does a data plane become an "OS done right"?

Edouard Bugnion (EPFL Lausanne, CH)

Main reference G. Prekas, M. Primorac, A. Belay, C. Kozyrakis, E. Bugnion, "Energy proportionality and workload consolidation for latency-critical applications," in Proc. of the 6th ACM Symp. on Cloud Computing (SoCC'15), pp. 342–355, ACM, 2015.
URL http://dx.doi.org/10.1145/2806777.2806848
Main reference A. Belay, G. Prekas, A. Klimovic, S. Grossman, C. Kozyrakis, E. Bugnion, "IX: A Protected Dataplane Operating System for High Throughput and Low Latency," in Proc. of the 11th USENIX Symp. on Operating Systems Design and Implementation (OSDI'14), pp. 49–65, USENIX Association, 2014.
URL https://www.usenix.org/conference/osdi14/technical-sessions/presentation/belay

In a (software) data plane architecture, the application bypasses the OS for most interactions and in particular I/O, e.g. using DPDK or some comparable user level networking framework. Data plane increase throughput and reduce jitter by first statically provisioning resources (cores and network queues) and then combining networking polling, bounded batching, and run-to-completion techniques. They can be build within an application [4, 5], by multi-threading networking with application [3], on top of research operating systems [1], or using hardware virtualization [2].

Of course, the focus on certain metrics comes at the expense of others, and points to two key limitations not found in conventional operating systems. (1) The static allocation of core that poll network queues leads to poor resource efficiency, e.g. energy proportionality or workload consolidation; (2) the run-to-completion model leads to a FCFC scheduling disciplines, which works fine in some cases (e.g., when the service time of packets is centrally distributed), but not for long-tailed or bi-modal situations.

The open question is whether #1 and #2 can be addressed without reducing the performance of the data plane model, or whether the feature creep will turn the data plane model into a regular operating system architecture. My immediate published work addresses problem #1 through a dynamic resource controls and low-level flow group migration mechanisms. Our current (very raw) research agenda looks at problem #2 through a data plane scheduler that combines FCFS with processor sharing.

References

1. Simon Peter, Jialin Li, Irene Zhang, Dan R. K. Ports, Doug Woos, Arvind Krishnamurthy, Thomas Anderson, and Timothy Roscoe. *Arrakis: The Operating System is the Control Plane*. OSDI, 2014.
2. Adam Belay, George Prekas, Ana Klimovic, Samuel Grossman, Christos Kozyrakis, and Edouard Bugnion. *IX: A Protected Dataplane Operating System for High Throughput and Low Latency*. OSDI, 2014.
3. EunYoung Jeong, Shinae Wood, Muhammad Jamshed, Haewon Jeong, Sunghwan Ihm, Dongsu Han, and KyoungSoo Park. *mTCP: a Highly Scalable User-level TCP Stack for Multicore Systems*. NSDI, 2014.
4. Hyeontaek Lim, Dongsu Han, David G. Andersen, and Michael Kaminsky. *MICA: A Holistic Approach to Fast In-Memory Key-Value Storage*. NSDI, 2014.
5. Ilias Marinos, Robert Watson, and Mark Handley. *Network stack specialization for performance*. ACM SIGCOMM, 2014.

3.3 Consensus as a Network Service

Marco Canini (University of Louvain, BE)

Joint work of Marco Canini, Tu Dang, Pietro Bressana, Han Wang, Ki Suh Lee, Hakim Weatherspoon, Fernando Pedone, Robert Soulé

The Paxos protocol is the foundation for building many fault-tolerant distributed systems and services. This talk posits that there are significant performance benefits to be gained by implementing Paxos logic in network devices. Until recently, the notion of a switch-based implementation of Paxos would be a daydream. However, new flexible hardware and expressive data plane programming languages are on the horizon and will provide customizable packet processing pipelines needed to implement Paxos. This talk describes an implementation of Paxos in one of those languages, P4, as well as our on-going efforts to evaluate the implementation on a variety of hardware devices. Implementing Paxos provides a critical use case for P4, and will help drive the requirements for data plane languages in general. In the long term, we imagine that consensus could someday be offered as a network service, just as point-to-point communication is provided today.

References

1 Huynh Tu Dang, Daniele Sciascia, Marco Canini, Fernando Pedone, and Robert Soulé. *NetPaxos: Consensus at Network Speed*. SOSR, 2015.

2 Huynh Tu Dang, Marco Canini, Fernando Pedone, and Robert Soulé. *Paxos Made Switch-y*. SIGCOMM Comput. Commun. Rev., Apr., 2016.

3 Huynh Tu Dang, Pietro Bressana, Han Wang, Ki Suh Lee, Hakim Weatherspoon, Marco Canini, Fernando Pedone, and Robert Soulé. *Network Hardware-Accelerated Consensus*. abs/1305.7429, 2016.

3.4 Recent Advances in Coflow-Based Networking

Mosharaf Chowdhury (University of Michigan – Ann Arbor, US)

Joint work of Mosharaf Chowdhury, Ion Stoica, Yuan Zhong, Zhenhua Liu, Matei Zaharia
Main reference M. Chowdhury, Y. Zhong, I. Stoica, "Efficient coflow scheduling with Varys," in Proc. of the 2014 ACM Conf. of the Special Interest Group on Data Communication (SIGCOMM'14), pp. 443–454, ACM, 2014.
URL http://dx.doi.org/10.1145/2619239.2626315

Over the past decade, the confluence of an unprecedented growth in data volumes and the rapid rise of cloud computing has fundamentally transformed systems software and corresponding infrastructure. To deal with massive datasets, more and more applications today are scaling out to large datacenters. Communication between the distributed computation tasks of these applications often result in massive data transfers over the network. Consequently, concentrated efforts in both industry and academia have gone into building high-capacity, low-latency datacenter networks at scale. At the same time, researchers and practitioners have proposed a wide variety of solutions to minimize flow completion times or to ensure per-flow fairness based on the point-to-point flow abstraction that forms the basis of the TCP/IP stack.

We observe that despite rapid innovations in both applications and infrastructure, application- and network-level goals are moving further apart. Data-parallel applications care about all their flows, but today's networks treat each point-to-point flow independently. This fundamental mismatch has resulted in complex point solutions for application developers, a myriad of configuration options for end users, and an overall loss of performance.

The recently proposed coflow abstraction bridges the gap between application-level performance and network-level optimizations. Each multipoint-to-multipoint coflow represents a collection of flows with a common application-level performance objective, enabling application-aware decision making in the network. We describe complete solutions including architectures, algorithms, and implementations that apply coflows to multiple scenarios using central coordination, and we demonstrate through large-scale cloud deployments and trace-driven simulations that simply knowing how flows relate to each other is enough for better network scheduling, meeting more deadlines, and providing higher performance isolation than what is otherwise possible using today's application-agnostic solutions.

3.5 A New Software Stack for Low-Latency Datacenters

John Ousterhout (Stanford University, US)

Hardware latencies for communication and storage access in datacenters are dropping dramatically, from milliseconds to (soon) microseconds. Unfortunately, today's highly-layered software stacks have too much overhead to capitalize on these improvements. If applications are to take advantage of the hardware improvements, we will need to make radical changes to the software stack. In this talk I will describe two projects underway at Stanford to create elements of a low-latency software stack. The first project is defining a new transport protocol for low-latency remote procedure calls; the second project is creating a new core-aware thread scheduling mechanism that moves most of the scheduler to user-level. This work, and other recent developments such as the rise of virtual machines, suggest that operating systems of the future may be "hollowed out": the data plane will largely bypass the operating system, leaving it mostly as a control plane.

3.6 NDP: New Datacenter Philosophy/Protocol

Costin Raiciu (University Politehnica of Bucharest, RO)

Joint work of Mark Handley, Alexandru Agache, Marcin Wojcik, Gianni Antiki, Costin Raiciu

Modern datacenter networks provide very high capacity and low switch latency, but transport protocols rarely manage to deliver performance matching the underlying hardware. We present NDP, a novel datacenter transport architecture that achieves both near-optimal file completion times and throughput in a wide range of scenarios including incast. NDP achieves this though a novel and synergisic combination of techniques including packet spraying, extremely small switch buffers, a randomized variant of packet trimming when queues fill, a priority queuing mechanism and a highly tuned receiver-driven ack clocking mechanism that can trigger retransmissions so rapidly that lost packets have very little impact on performance.

3.7 Where has the time gone?

Noa Zilberman (University of Cambridge, GB)

Joint work of Matthew Grosvenor, Neelakandan Manihatty-Bojan, Diana Andreea Popescu, Gianni Antichi, Salvator Galea, Andrew Moore, Robert Watson, Marcin Wojcik, Noa Zilberman

Latency control is mandatory to the operation of data centres. In this talk we present a study of the different latency components and show the breakdown of end-to-end latency, all the way from the application level to the wire. This breakdown provides insights on the challenges for latency control and the importance of different latency components. We present several trajectories taken in Cambridge to improve latency control, and discuss the instrumentation required to this end.

4 Working groups

4.1 Low latency outside the DC (5G, open Internet, etc)

Jari Arkko (Ericsson – Jorvas, FI), Bob Briscoe (Simula Research Laboratory – Lysaker, NO), Jon Crowcroft (University of Cambridge, GB), Koen De Schepper (NOKIA Bell Labs – Antwerp, BE), Lars Eggert (NetApp Deutschland GmbH – Kirchheim, DE), Michio Honda (NetApp Deutschland GmbH – Kirchheim, DE), Wolfgang Kellerer (TU München, DE), Kirill Kogan (IMDEA Networks – Madrid, ES), Mirja Kühlewind (ETH Zürich, CH), and Michael Scharf (NOKIA – Stuttgart, DE)

Discussion points:

- Clarification: use case providers need to clearly distinguish between guaranteed latency and low latency.
- Localized computing applications can potentially benefit from low-latency.
- 5G and DC networking have common aspects and can be considered similar in some respects: routing and addressing are challenges addressed by both, but do not scale to both cases. Capacity is a key concern in both cases but has fundamental differences: capacity comes from multi-paths in DC and frequency sharing in 5G. Can lessons learned in one domain be carried over to another? E.g., can multi-path routing be carried over to 5G to optimize the capacity problem?
- Limitations: Internet is a signaling plane. DC mechanisms can not be deployed on Internet.
- There is no DC position. In IETF, DC things do not run in Routing area, DC proposals got pushed to transport. Perhaps change IETF to take in IP layer work.
  - routing area in IETF, transport area more focused
  - Specification is not regarded as publication, a motivation issue for academics
- Internet needs a framework to bring in DC technologies.

4.2 How to make datacenter research more reproducible

Marinho Barcellos (Federal University of Rio Grande do Sul, BR), Olivier Bonaventure (University of Louvain, BE), Edouard Bugnion (EPFL Lausanne, CH), Oliver Hohlfeld (RWTH Aachen, DE), Andrew W. Moore (University of Cambridge, GB), Hakim Weatherspoon (Cornell University – Ithaca, US), and Noa Zilberman (University of Cambridge, GB)

Question: How do we make artifacts reproducible in the context of data centers?

Summary/Takeaways:

1. whatever we do have to have incentives and has to scale
2. Back end: longer repeatable papers journal; specifically CCR
3. Front end: Track for reproducibility at SIGCOMM
4. Community award for reproducibility
5. Badging: Use ACM policy (see Results and Artifact Review and Badgin – http://www.acm.org/publications/policies/artifact-review-badging/

4.3 Switch programmability: How much functionality should be handled by switches?

Matthew P. Grosvenor (University of Cambridge, GB)

- Starting point: Innovators dilemma – building components means that you already have a model of how the components might go together which means they need to envisage how other innovators might innovate creates constrictions on what they can do. Should really only do things in HW that become stable over de
- Current SDN concentrates on efficient representation of packet processing, but not buffer management. Desire to make buffer/queue management programmable.
- P4 building blocks: P4 pipeline, match packet → place into buffer → (drop, send, modify) P4 may be able represent some states, but may not be able to implement these states efficiently? If we consider a single queue, there are only two things we are doing: ingress and egress and we want to be able to program that.
- But we may want to do more than just simple "old" "boring" "classic" "legacy" switch operations. We want to do packet rewriting, paxos, content based load balancing etc. etc. What are the primitives to do this. For NetPaxos primitives: state/register, compare, broadcast.
- What are the sorts of applications outside of networks that we can use to implement building blocks. E.g. stream processing, aggregation etc. that we can use to steer design decisions.
- Cisco what are the applications? We want to use from the switches? What sort of functionality do we need? What sort of monitoring? Statistics? Average queuing delay. Monitoring may be at a huge cost e.g. 25,000 queues counters, costs energy etc.
- Straw man: In a datacenter we only need two things: source routing and priority queueing Counter argument: You need some kind of feedback from the network (not packet drop?) not ECN? But what? Perhaps packet.
- NIC features / Switch and NIC vs. Switch Are they the same? The NIC may want to offload more complex things into the CPU Or maybe the division of labour between NIC and switch.
- How are NICs different to switches. Application offload (TCP segmentation), encryption, control groups NICs have a more information. More generic building blocks necessary in the NIC for other protocols. Experience in NICs suggests that NICs are slow, and offloading into the NIC from a high speed CPU usually has no net positive gain.
- What about C# → gates? Is that enough? Application programmer wants to know what the template engine is that does C# → gates? E.g. template languages that were used in GPUs. Do we really want the application program to run on the NIC, or do we want some generic functions in the NIC, that many applications do want to do? What about application sharing? How do we share the resources in the NIC.

4.4 Congestion Control for 100 G networks

Dongsu Han (KAIST – Daejeon, KR), Alexandru Agache (University Politehnica of Bucharest, RO), Mohammad Alizadeh Attar (MIT – Cambridge, US), Marco Canini (University of Louvain, BE), Nandita Dukkipati (Google Inc. – Mountain View, US), Matthew P. Grosvenor (University of Cambridge, GB), Patrick Jahnke (TU Darmstadt, DE), Lavanya Jose (Stanford University, US), Mike Marty (Google – Madison, US), John Ousterhout (Stanford University, US), Rong Pan (CISCO Systems – San Jose, US), Ihsan Ayyub Qazi (LUMS – Lahore, PK), Costin Raiciu (University Politehnica of Bucharest, RO), and Michael Scharf (NOKIA – Stuttgart, DE)

Topic: What is different in 100 Gbps?

1. Everything may look like mice?
   - Argument: everything look like mice. ↔ counter: some apps will find way to fill up.
   - CPU (a single thread, single flow) may not be that fast for a single flow too fill up 100 Gbps ↔ RDMA (cpu offload) is there to fill up.
   - Today's traffic distribution comes from today's limitation. If b/w is plenty, some apps will find way to fill the pipe.
     - Everything looking like a mice → What if they arrive altogether?
     - Solving incast very hard at 100 Gbps. (no time to control)
     - Counter argument: But because it requires tighter synchronization at 100 G, it may not matter if there enough randomness in the host (assuming flow size is the same).
     - Even a small randomness at the sender will break synchronization from the network standpoint.
     - Counter argument: tomorrow flow size will increase (as b/w will become plenty)
   - Conclusion: We bet on that there will be apps that send more flows. If we have short incast, congestion control can't do much in at 100 G. Or we don't have problem b/c of plenty b/w and less synchronized flows.
2. Core oversubscription
   - Even today, core of data center is oversubscribed due to cost. Full bisection B/W is a myth. With 100 G at the host, oversubscription will become worse because it so expensive to provision for full bisection.
   - Note, 1G → 10G transition was the opposite.
   - Potentially, this will increase core congestion.
3. Buffers
   - Buffers are getting scares per port per gbps. Sometimes even absolute amount of buffers is going down.
   - How do you cope with incast, large fan-in without large buffers?
   - Some high-end 100 G switches have no shared buffer.
   - Packets are scattered throughout non-shared buffers. Difficult to count the number of packets, which is required for ECN. Timing also difficult. So delay feedback won't work.
   - Congestion control alternatives:
     - Loss as congestion signal always our friend.

- If you have very low buffers, credit-based flow control looks more attractive, which means IB.
- On-chip networks do credit-based flow control.

4.5 Role of the Operating System (Group 2)

Mike Marty (Google – Madison, US), Jari Arkko (Ericsson – Jorvas, FI), Patrick Thomas Eugster (TU Darmstadt, DE), Oliver Hohlfeld (RWTH Aachen, DE), and Andrew W. Moore (University of Cambridge, GB)

Slogan: OS job is to allocate, map, and multiplex if necessary.

Problems we have right now:

- low productivity dev environment (old time shared machine perception leads to the current OS concept)
- control plane done wrong or data plane done right?
- hard to test and debug if there are no other ways to do isolation
- unclear picture: API? or privileges?
- multi-core scheduling (e.g., Linux)
- OS is naive in resource sharing beyond CPU and process management, particular for networking – hard to evolve the OS

Discussion points:

- Is the core OS role to aid the development. or then aid for debugging?
- OS is good to do fine-grain sharing
- If hardware can multiple, then OS only does the map, then scheduling becomes a allocation.
- DC OS and stand along OS?
- Accounting and auditing role for OS?
- Open questions: applications may need know the need, but express such needs in a less-defined abstraction way. The needs from application should be abstracted, otherwise people will do things for each of them

4.6 Role of the Operating System (Group 1)

John Ousterhout (Stanford University, US), Edouard Bugnion (EPFL Lausanne, CH), Marco Canini (University of Louvain, BE), Nandita Dukkipati (Google Inc. – Mountain View, US), and Jitendra Padhye (Microsoft Corporation – Redmond, US)

- Docker style containers are popular but inter-container communication is not the same as IPC. Can we have isolation features without the communication overhead of going through all the layers to communicate b/n two containers?
  Unikernels try to solve the problem by compiling a custom OS for the application at runtime with just the necessary layers.

- Why is kernel networking so slow?
  - Guest VMs in clouds have to go through two network stacks. Hypervisor bypass at least allows them to skip the hypervisor and (think they can) run directly on NICs. Only Amazon has done this successfully- maybe using bump in the wire (i.e., CPU on other side of NIC modifying packets at line rate...)
  - Most of the time is spent in the socket layer, not the TCP/IP processing?
  - Sockets as an API b/n the application and networking stack doesn't lend itself to scaling with multicores. There is typically a lot if churn in sockets. In fact Google's SPDY was just a way for user to open socket once and have a permanent connection to the mothership. Sockets is the wrong API.
- What is the right API then? Maybe RDMA?
  - RDMA is better than RPC at the tail (it's one sided, while RPC could suffer cu of interrupts at destination etc.) Heated discussion on whether RDMA was better for median. J.O.: RDMA memory model bad for complex applications, if multiple sources want to update a key, value they'll have to get locks, look up data structures to find new memory for a bigger value etc..
- Rack scale computing vs tighter NIC-CPU integration for super-fast networks
  - No interesting application that can fit in a rack yet but hardware will evolve. Why do we need different interconnects within a rack? Can go CPU to CPU, avoiding PHY, learning fabric etc. If it's latency may be we can just avoid the half us PCIe latency with having the NIC right next to the fastest cache.
- Hardware offloading gives some performance saving and orders of magnitude power savings compared to function virtualization, is the way to go. What about ossification? Re-configurable smart NICs that can be programmed using P4 enable new features to be offloaded (yet to characterize the kind of features).

Suggested topics for an extended discussion:

- Thread scheduling
- Role of the NIC and the OS
- Storage class memory / NVRAM
- Impact of Isolation on Performance: e.g., Containers are Processes Isolated form each other, when they want to talk to each other through network., overhead
- What is the right API between applications and OS bypass?
- OSes and distributed application, OS : what's the boundary between what OS schedules for one machine and what scheduler schedules for a cluster of machines
- Disaggregated computing: Disaggregation enabled by low latency interconnects "Rack scale computing"

5 Panel discussions

5.1 Concluding Statements

John Ousterhout (Stanford University, US)

In the Wednesday morning session, seminar participants proposed possible conclusions from our discussions. Each conclusion was discussed briefly, then voted on. The proposed conclusions and votes follow below.

- 5–15 us µs round-trip latencies will soon be common in datacenter applications
  Yes: 18 , No: 3
  - Depends if we localize computation. In Google DCs, the speed of light is higher than that (100's of meters fibers). Todo: do a better job in taking hierarchy into account
  - Every DC is different. Googles DC is just orders of magnitude larger than others.
- Packet scheduling (clock-out) should be done by the receiver (or its TOR), as much as possible
  Yes: 9, No: 15
- Network oversubscription will always be present in the datacenter (apps will be bandwidth-constrained by core bandwidth)
  Yes: 20, No: 5
  - At what time-scale? You can't assume that you'll never run into congestion etc.
  - Question is more on which apps are bandwidth constrained or not?
- Future datacenters will be flat (not pod-structured)
  Yes: 3.5, No: 16
  - Key observation: if we would have this seminar 4 years ago, we wouldn't have this discussion. The common thinking at the time was to take layer 3 all the way down. Is this just an artifact of current thinking?
- Link-level flow control will become a primary solution at datacenter scale (i.e. no congestion drops)
  Yes: 5, No: 14
  - Long-term: source aggressive senders down – could not become a solution. Short-term is possible by link-level flow control
  - Building a zero loss networks is expensive
  - Remind lessons learned on ATM control loops failed: ABR (available bit rate) operated edge to edge across the VC through all the switches- problem was that the rate adjustment cycle was close to TCP's AIMD scheme, but not quite, so the two, nested control loops of TCP & ATM would be out of phase, and so the capacity/demand would oscillate, between underutilized, and overusing/ with higher delay. Because there was an IP layer between TCP and the ATM VC, there wasn't an easy way to get cross layer info from the lower layer up to TCP (always assuming we could have had a way to inform TCPs cwnd adjustment algorithm with something more sophisticated than just delay or loss or ECN, which we didn't have).

The following conclusions were voted on with no discussion:

- The OS should not be involved in in-host dataplane networking ops (read and write calls)
  Yes: 18, No: 3
- The OS should allocate, map, and multiplex only when necessary
  Yes: 21, No: 0
- DC and rest of the Internet are different (speed-of-light)
  Yes: 18, No: 4
- Link-level flow control will become a primary solution at rack scale
  Yes: 16, No: 4

6 Open problems

6.1 Questions about flow aggregation

Michael Welzl (University of Oslo, NO)

In a practical datacenter setting, how can we apply flow (actually, congestion control) aggregation in data centers? By doing congestion control between hypervisors? Would this have to be carrying TCP traffic, or can we develop our own specialized congestion control? Can we aggregate at TOR switches?

Participants

- Alexandru Agache
University Politehnica of Bucharest, RO
- Mohammad Alizadeh Attar
MIT – Cambridge, US
- Jari Arkko
Ericsson – Jorvas, FI
- Marinho Barcellos
Federal University of Rio Grande do Sul, BR
- Olivier Bonaventure
University of Louvain, BE
- Bob Briscoe
Simula Research Laboratory – Lysaker, NO
- Edouard Bugnion
EPFL Lausanne, CH
- Marco Canini
University of Louvain, BE
- Julian Chesterfield
University of Cambridge, GB
- Mosharaf Chowdhury
University of Michigan – Ann Arbor, US
- Jon Crowcroft
University of Cambridge, GB
- Koen De Schepper
NOKIA Bell Labs – Antwerp, BE
- Aaron Yi Ding
TU München, DE
- Nandita Dukkipati
Google Inc. – Mountain View, US
- Lars Eggert
NetApp Deutschland GmbH – Kirchheim, DE
- Patrick Thomas Eugster
TU Darmstadt, DE
- Matthew P. Grosvenor
University of Cambridge, GB
- Dongsu Han
KAIST – Daejeon, KR
- Oliver Hohlfeld
RWTH Aachen, DE
- Michio Honda
NetApp Deutschland GmbH – Kirchheim, DE
- Patrick Jahnke
TU Darmstadt, DE
- Lavanya Jose
Stanford University, US
- Wolfgang Kellerer
TU München, DE
- Kirill Kogan
IMDEA Networks – Madrid, ES
- Mirja Kühlewind
ETH Zürich, CH
- Mike Marty
Google – Madison, US
- Andrew W. Moore
University of Cambridge, GB
- John Ousterhout
Stanford University, US
- Jitendra Padhye
Microsoft Corporation – Redmond, US
- Rong Pan
CISCO Systems – San Jose, US
- Max Plauth
Hasso-Plattner-Institut – Potsdam, DE
- Ihsan Ayyub Qazi
LUMS – Lahore, PK
- Costin Raiciu
University Politehnica of Bucharest, RO
- Michael Scharf
NOKIA – Stuttgart, DE
- Hakim Weatherspoon
Cornell University – Ithaca, US
- Klaus Wehrle
RWTH Aachen, DE
- Michael Welzl
University of Oslo, NO
- Noa Zilberman
University of Cambridge, GB

Report from Dagstuhl Seminar 16282

Topological Methods in Distributed Computing

Edited by

Dmitry Feichtner-Kozlov[1] and Damien Imbs[2]

1 **Universität Bremen, DE,** dfk@math.uni-bremen.de
2 **Universität Bremen, DE,** imbs@math.uni-bremen.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 16282 "Topological Methods in Distributed Computing", which was attended by 22 international researchers, both junior and senior. In the last 10–15 years, there has been an increasing body of work on applications of topology in theoretical distributed computing. This seminar brought together computer scientists working in theoretical distributed computing and mathematicians working in combinatorial topology, leading to interesting new collaborations. This report gathers abstracts of the talks given by the participants and of the group research sessions that happened during the seminar.

Seminar July 10–15, 2016 – http://www.dagstuhl.de/16282
1998 ACM Subject Classification C.2.4 Distributed Systems, H.1.1 Systems and Information Theory
Keywords and phrases combinatorial topology, concurrency, distributed protocols, shared-memory communication, solvability
Digital Object Identifier 10.4230/DagRep.6.7.31

1 Executive Summary

Dmitry Feichtner-Kozlov
Damien Imbs

This seminar brought together 22 researchers in combinatorial topology and in theoretical distributed computing. Participants came from Germany, France, Israel, Switzerland, United States, Canada, Japan and Austria. The seminar featured a combination of 1-hour talks, group sessions and an open problems session.

Scientific background and topics of the seminar

In the classical sequential computational model, computability is viewed through the Church-Turing thesis, where computations are reduced to those done by Turing machines, and complexity issues are of central importance. In the distributed setting, the situation is quite different. Since the threads of executions may intertwine in various ways (depending on the model), one of the central issues becomes dealing with execution ambiguity, and deciding whether certain standard tasks (Consensus, Renaming, etc.) are computable at all.

In this sense, the distributed setting is harder to analyze rigorously than the sequential one, or at least the difficulties are of quite different type. At the same time very many

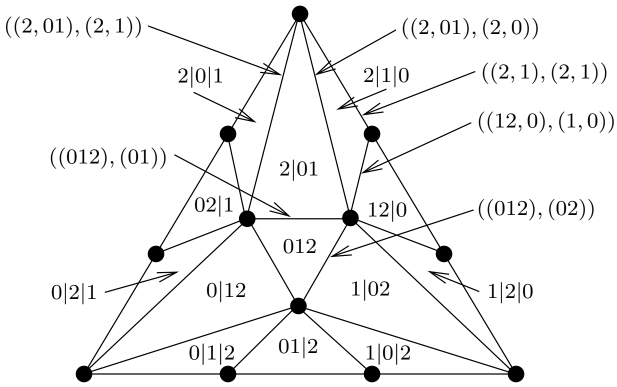

Figure 1 The protocol complex for the one-round execution of the standard immediate snapshot protocol for 3 processors.

real-life situations need to be modeled by the distributed setting. These include networks of banking machines, or networks of flight controllers and airplanes, who need to reach a common decision in a decentralized setting. Another example is the parallel chip design, where we need to understand what type of elementary operations – so-called computational primitives, have to be implemented on the hardware level, so that the resulting computational system is powerful enough for our needs.

In the 80s it was realized (due in particular to the work of Fischer, Lynch and Paterson) that certain standard tasks (Consensus) cannot be solved in standard computational models (such as Message-Passing) in the presence of even simple processor crash failures. As spectacular as it is, it has become one of the steps in the development of a sophisticated and beautiful subject of theoretical distributed computing; we refer here to the classical books of Lynch and Attiya & Welch.

In the late 90s and in the early years of our millenium, it was realized by at least 3 independent groups of researchers that topological methods are applicable in proving impossibility results in theoretical distributed computing. There followed a process of further penetration of topological methods, which by now have gained a definite foothold in distributed computing. Additionally, there has also been some work on mathematical foundations, though much remains to be done when it comes to precise definitions and rigorous proofs. Independently, we feel that it is of great interest to develop the mathematics which is inspired by these methods.

The state-of-the-art of the subject was recently summarized in a book by Herlihy, Kozlov and Rajsbaum. One of the paradigms introduced there is to replace the computational task specification by the triple: input complex, output complex, and task specification map, there the input and the output complexes are simplicial complexes with additional structure, and the task specification map is what we call a carrier map, whose definition reflects our desire to restrict ourselves to the wait-free protocols. All the wait-free tasks can be encoded this way, and as a result one obtains both well-known as well as new structures from combinatorial topology.

Furthermore, one can consider the simplicial model for the totality of all executions of a given protocol – the so-called protocol complex. In the full formal setting one actually considers a triple of two simplicial complexes and a carrier map, each one equipped with an additional structure. The intuition here is that the second simplicial complex, as well as the carrier map depend heavily on the model of computation that we choose. One standard example is to take the so-called Immediate Snapshot model. On Figure 1 we show the protocol complex for the one-round execution of the standard immediate snapshot protocol for 3 processors. As already this example shows, frequently there is a purely combinatorial description of the arising simplicial structure. The question of wait-free computability of a given task in a given computational model reduces then to the question of existence of a simplicial map from the protocol complex to the output complex, the so-called decision map, which satisfies certain conditions, which in essence mean that the outputs obtained by the protocol are valid under the task specification. Furthermore, we also have mathematical models for anonymous tasks, and anonymous protocols, as well as for colorless tasks.

As one can see, the mathematics needed for the current model is essentially that of simplicial complexes and carrier maps between them. With subsequent deepening of the theory and diversification of the considered questions, many further mathematical fields are coming in: for example, one needs to consider group actions and equivariant maps, as well as simplicial and carrier maps which satisfy other, less standard conditions. Many of the questions which arise in this setup are somewhat different from the questions classically studied in the simplicial context.

2 Table of Contents

3 Overview of Talks

3.1 The Entropy of Shared Memory

Peva Blanchard (EPFL – Lausanne, CH) and Julien Stainer (EPFL – Lausanne, CH)

We ask to what extent two processes communicating through shared memory can extract randomness from a stochastic scheduler, e.g., to generate random numbers for cryptographic applications. We introduce the quantitative notions of entropy rate and information capacity of a distributed algorithm. Whilst the entropy rate measures the Shannon information that may pass from a given scheduler to the processes executing the algorithm, the information capacity measures the optimal entropy rate over all possible schedulers.

We present a general method for computing these quantities by classifying distributed algorithms according to their pattern of shared memory accesses. We then address the issue of effectively extracting, online, this information into a meaningful format at every process. We present an algorithm solving this problem with an optimal memory consumption.

3.2 Anonymous Graph Exploration with Binoculars

Jeremie Chalopin (Aix-Marseille University, FR)

Joint work of Jeremie Chalopin, Emmanuel Godard, Antoine Naudin
Main reference J. Chalopin, E. Godard, A. Naudin, "Anonymous Graph Exploration with Binoculars", in Proc. of the 29th Int'l Symp. on Distributed Computing (DISC'15), LNCS, Vol. 9363, pp. 107–122, Springer, 2015.
URL http://dx.doi.org/10.1007/978-3-662-48653-5_8

We investigate the exploration of networks by a mobile agent. It is long known that, without global information about the graph, it is not possible to make the agent halts after the exploration except if the graph is a tree. We therefore endow the agent with binoculars, a sensing device that enables the agent to see the graph induced by its neighbors.

We show that, with binoculars, it is possible to explore and halt in a large class of non-tree networks. We give a complete characterization of the class of networks that can be explored using binoculars using standard notions of discrete topology. This class is much larger than the class of trees: it contains in particular chordal graphs, plane triangulations and triangulations of the projective plane. Our characterization is constructive: we present an Exploration algorithm that is universal; this algorithm explores any network explorable with binoculars, and never halts in non-explorable networks.

3.3 Introduction to matroids (with a view on computations in topology)

Emanuele Delucchi (University of Fribourg, CH)

Matroid theory is a structurally rich branch of combinatorics with pervasive connections with geometry, topology, and algebra.

In this expository talk I will gently introduce the notion of matroid and illustrate some of its applications and main open problems. Then – with an view towards topological applications – I will go on to more refined combinatorial structures such as oriented matroids. I will conclude with a recent and already thriving new character among matroidal theories: Baker's matroids over hyperfields.

3.4 Evasiveness

Etienne Fieux (Université Paul Sabatier – Toulouse, FR)

This is an introductory talk on evasiveness. It will not give a presentation of recent results on this question, the aim being to give an illustration of the use of tools from algebraic topology in order to solve questions in combinatorics. More precisely, we will emphasize the role of collapsibility for making a bridge between the original version of the evasiveness conjecture (about query complexity of boolean functions) and the "topological" version in terms of simplicial complexes. We also will present important results obtained by this way.

3.5 What is Distributed Computability? (Personal Perspective)

Eli Gafni (UCLA, US)

Early research in Distributed-Computing (DC) à la Dijkstra and Lamport centered on the Mutual-Exclusion (MX) problem. In this problem a processor with distinguished privilege, has to signal to other processors that it relinquishes its privilege. Hence, the "Wait-until Condition" construct naturally had to be used. They did not stop and ask whether it is essential in a one-shot case of just securing the privilege once.

In a landmark paper that started distributed-computability, Fischer, Lynch and Patterson in 1983 observed that the one-shot model of MX is of runs such that any processor that took a single step in a run will continue to take steps infinitely-often. They then showed that any meaningful relaxation of this set of runs in the form of allowing a single processor out of any number, to appear only finite number of times greater than 0, is an obstruction to the ability to secure the privilege. Thus, the power of coordination depends on the set of runs under consideration: Coordinating between processors that have limited agreement about what run they might be in. Thus, DC is the study of the coordination power of a given set of runs. Theoreticians found it convenient to consider models which are not "linear" like runs. Nevertheless, such models should be used only insofar as they correspond to some set of runs, i.e. the coordination-power of the non-linear model under study should be exactly the same as the coordination power of a set of runs.

3.6 Topological methods for distributed computability in communication networks

Emmanuel Godard (Aix-Marseille University, FR)

Coverings and fibrations are topological morphisms that capture very well the locality of distributed computations in communication networks. First introduced to generalize the seminal impossibility result of Angluin in anonymous networks, they proved to be a very versatile tool to get complete distributed computability characterizations in many models. We present here a general framework to investigate distributed computability in communication networks (that could be anonymous, homonymous or with ids) for any scheduler, both for deterministic and probabilistic algorithms. This framework comprehends all known characterisations and provides a direct and effective methodology to investigate new models.

3.7 Homotopy theory and concurrency

John Frederick Jardine (University of Western Ontario – London, CA)

Higher dimensional automata are defined as finite cubical complexes. The path category P(K) of such an object K has objects all vertices (states) of K, and its morphisms are the execution paths between states. An algorithm for computing P(K) is derived from an explicit 2-category resolution. Complexity reduction methods are discussed. A first parallel technique for computing P(K) is displayed.

3.8 Concurrency as an Iterated Affine Task

Petr Kuznetsov (Telecom Paris Tech, FR)

We consider models of computations expressed via sets of runs bounding the concurrency level: the number of processes that can be concurrently active. The model of k-concurrency is equivalent to the read-write shared-memory systems equipped with k-set-agreement objects. We show that every such model can be characterized by an affine task: a simple combinatorial structure (a simplicial complex), defined as a subset of simplices in the second degree of the standard chromatic subdivision. Our result implies the first combinatorial representation of models equipped with abstractions other than read-write registers.

3.9 Tight Bounds for Connectivity and Set Agreement in Byzantine Synchronous Systems

Hammurabi Mendes (University of Rochester, US)

Joint work of Hammurabi Mendes, Maurice Herlihy

In this talk, we discuss how the protocol complex of a Byzantine synchronous system can remain $(k-1)$-connected for up to $\rceil t/k \lceil$ rounds, where t is the maximum number of Byzantine processes, with $t \geq k \geq 1$. In crash-failure systems, the connectivity upper bound is $\rfloor t/k \lfloor$ rounds, so Byzantine ambiguity can potentially persist for one extra round, delaying the solution to k-set agreement and other related problems. We see how our connectivity bound is tight, at least when the number of processes is suitably large compared to the number of failures, by solving an appropriate formulation of k-set agreement in $\rceil t/k \lceil +1$ rounds.

3.10 Carrier Complex: A Poset Topology for Finding Distributed Protocols

Susumu Nishimura (Kyoto University, JP)

We study the problem of finding a wait-free protocol for read-write shared memory distributed systems that meets a given topological specification of colored distributed task. Due to the asynchronous computability theorem by Herlihy and Shavit, this amounts to finding a pair of a subdivision on the input complex and a chromatic simplicial map onto the output complex, which is an undecidable problem.

Though there can be no perfect solution to the problem, we claim that there is a partial solution. Given a carrier map as a task specification, we first set up a variant of a well-known poset topology, called the carrier complex, which is a simplicial complex obtained by a combinatorial construction on the carrier map. The solution is then reached by finding an appropriate sequence of simplicial collapses that deforms the input complex embedded as a subcomplex within the carrier complex.

To put the above method to work for concrete instances of tasks in practice, we have to find a reasonably efficient algorithm. The complexity bound of the algorithm is polynomial when the number of processes is fixed, but the constant factor rapidly increases for a larger number of processes. Based on a preliminary experiment on a prototype implementation, we report that the algorithm is expected to work in practice for up to 4 processes.

3.11 Algebraic Methods in Distributed Graph Algorithms

Ami Paz (Technion – Haifa, IL)

Joint work of Keren Censor-Hillel, Petteri Kaski, Janne H. Korhonen, Christoph Lenzen, Ami Paz, Jukka Suomela
Main reference K. Censor-Hillel, P. Kaski, J. H. Korhonen, C. Lenzen, A. Paz, J. Suomela, "Algebraic Methods in the Congested Clique", in Proc. of the 2015 ACM Symp. on Principles of Distributed Computing (PODC'15), pp. 143–152, ACM, 2015.
URL http://dx.doi.org/10.1145/2767386.2767414

We will survey algorithm and lower bound techniques in three distributed message passing models – local, congest and congested clique. Then, we will discuss our recent results [PODC'15] regarding these problems: implementing matrix multiplication in the congested clique model, and using it to solve various distance computation and subgraph detection problems in this model.

If time suffices , we will discuss some restricted lower bounds for these problems, and the failure to achieve general lower bounds in the congested clique model.

3.12 Topological characterisation of 2-set agreement in communication networks with omission faults

Eloi Perdereau (Aix-Marseille University, FR)

We consider an arbitrary communication network G where at most f messages can be lost at each round. We give a topological characterisation when $G = K_3$ and $G = P_3$. We then show it is possible to reduce any arbitrary graph G to one of this case.

3.13 Easy Impossibility Proofs for k-Set Agreement

Ulrich Schmid (TU Wien, AT)

Despite of being quite similar agreement problems, distributed consensus (1-set agreement) and general k-set agreement require surprisingly different techniques for proving impossibilities. In particular, the relatively simple bivalence arguments used in the impossibility proof for consensus in the presence of a single crash failure are superseded by a reasoning based on algebraic topology resp. variants of Sperner's Lemma in the case of k-set agreement for $f \geq k > 1$ crash failures. In this talk, we provide an overview of a generic theorem for proving the impossibility of k-set agreement in various message passing settings, which has been published at OPODIS'11 [1]. Our BRS Theorem is based on a reduction to the consensus impossibility in a certain subsystem resulting from a partitioning argument, which facilitates easy impossibility proofs for k-set agreement. Its broad applicability is demonstrated in several message passing models.

The relevance of the BRS Theorem for this Dagstuhl Seminar results from the fact that both the core idea (reduction to consensus) and the main precondition (partitioning) of the BRS Theorem have counterparts in the algebraic topology setting. Exploring the exact nature of these is not only interesting in itself, but may also lead to a better understanding of the actual relation of partitioning-vs. Sperner-Lemma-based impossibility results.

References

1 Martin Biely, Peter Robinson and Ulrich Schmid, "Easy Impossibility Proofs for k-Set Agreement in Message Passing Systems", in Proc. of the 15th Int'l Conf. on Principles of Distributed Systems (OPODIS'11), LNCS, Vol. 7109, pp. 299–312, Springer, 2011. DOI: 10.1007/978-3-642-25873-2_21

4 Working groups and the open problems session

Four group sessions, of two and a half hours each, were organized in the afternoons. These sessions served as a complement to the lectures and provided a less formal setting in which the researchers could exchange ideas and locate areas in which progress could be made in the near future. The sessions also served as an excellent opportunity to jump-start joint projects. Due to the extensive interest and the limited amount of time available to the organizers, the group sessions were scheduled in parallel and both tracks were well attended.

One of the highlights was an extensive session on higher-dimensional expansion (by Roy Meshulam), a modern topic bringing a new approach to the classical graph expanders. Due to the technical nature of the subject and its novelty to most participants, the working session provided a very fitting setting for the discussions. Another topic which was addressed was the possible connections of higher homotopy and the study of concurrency to the distributed setting (by Rick Jardine). This later theme has already brought some fruit, and there are upcoming workshops which will continue in this new research direction.

The open problems session was also attended by everyone and resulted in lively discussions of several central problems. Just to mention one: the round-complexity of the Weak Symmetry Breaking, and the intricate mathematics connected to it, was the source of interest and fascination. It was clear to the participants, that in this, and in many other instances, much remains to be understood by the community.

We feel that the sessions were a success and made use of the unique setting provided at the Dagstuhl Research Center, versus a usual conference venue.

Participants

- Peva Blanchard
EPFL – Lausanne, CH
- Jeremie Chalopin
Aix-Marseille University, FR
- Emanuele Delucchi
University of Fribourg, CH
- Dmitry Feichtner-Kozlov
Universität Bremen, DE
- Etienne Fieux
Université Paul Sabatier – Toulouse, FR
- Eli Gafni
UCLA – Los Angeles, US
- Emmanuel Godard
Aix-Marseille University, FR
- Hans Hagen
TU Kaiserslautern, DE
- Damien Imbs
Universität Bremen, DE
- John Frederick Jardine
University of Western Ontario – London, CA
- Petr Kuznetsov
Telecom Paris Tech, FR
- Jan-Philipp Litza
Universität Bremen, DE
- Hammurabi Mendes
University of Rochester, US
- Roy Meshulam
Technion – Haifa, IL
- Susumu Nishimura
Kyoto University, JP
- Ami Paz
Technion – Haifa, IL
- Eloi Perdereau
Aix-Marseille University, FR
- Thibault Rieutord
Telecom Paris Tech, FR
- Matthieu Roy
LAAS – Toulouse, FR
- Ulrich Schmid
TU Wien, AT
- Jan Felix Senge
Universität Bremen, DE
- Julien Stainer
EPFL – Lausanne, CH

Report from Dagstuhl Seminar 16291

Data, Responsibly

Edited by

Serge Abiteboul[1], Gerome Miklau[2], Julia Stoyanovich[3], and Gerhard Weikum[4]

1 **ENS – Cachan, FR,** serge.abiteboul@inria.fr
2 **University of Massachusetts – Amherst, US,** miklau@cs.umass.edu
3 **Drexel University – Philadelphia, US,** stoyanovich@drexel.edu
4 **MPI für Informatik – Saarbrücken, DE,** weikum@mpi-inf.mpg.de

Abstract

Big data technology promises to improve people's lives, accelerate scientific discovery and innovation, and bring about positive societal change. Yet, if not used responsibly, large-scale data analysis and data-driven algorithmic decision-making can increase economic inequality, affirm systemic bias, and even destabilize global markets.

While the potential benefits of data analysis techniques are well accepted, the importance of using them responsibly – that is, in accordance with ethical and moral norms, and with legal and policy considerations – is not yet part of the mainstream research agenda in computer science.

Dagstuhl Seminar "Data, Responsibly" brought together academic and industry researchers from several areas of computer science, including a broad representation of data management, but also data mining, security/privacy, and computer networks, as well as social sciences researchers, data journalists, and those active in government think-tanks and policy initiatives. The goals of the seminar were to assess the state of data analysis in terms of fairness, transparency and diversity, identify new research challenges, and derive an agenda for computer science research and education efforts in responsible data analysis and use. While the topic of the seminar is transdisciplinary in nature, an important goal of the seminar was to identify opportunities for high-impact contributions to this important emergent area specifically from the data management community.

Seminar July 17–22, 2016 – http://www.dagstuhl.de/16291

1998 ACM Subject Classification H.2.8 Database Applications, H.3.3 Information Search and Retrieval, K.4.1 Public Policy Issues, K.6.5 Security and Protection

Keywords and phrases Data responsibly, Big data, Machine bias, Data analysis, Data management, Data mining, Fairness, Diversity, Accountability, Transparency, Personal information management, Ethics, Responsible research, Responsible innovation, Data science education

Digital Object Identifier 10.4230/DagRep.6.7.42

Data, Responsibly, *Dagstuhl Reports*, Vol. 6, Issue 7, pp. 42–71

Editors: Serge Abiteboul, Gerome Miklau, Julia Stoyanovich, and Gerhard Weikum

DAGSTUHL REPORTS Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive summary

Serge Abiteboul
Gerome Miklau
Julia Stoyanovich
Gerhard Weikum

Our society is data-driven. Large scale data analysis, known as Big data, is distinctly present in the private lives of individuals, is a dominant force in commercial domains as varied as automatic manufacturing, e-commerce and personalized medicine, and assists in – or fully automates – decision making in the public and private sectors. Data-driven algorithms are used in criminal sentencing – ruling who goes free and who remains behind bars, in college admissions – granting or denying access to education, and in employment and credit decisions – offering or withholding economic opportunities.

The promise of Big data is to improve people's lives, accelerate scientific discovery and innovation, and enable broader participation. Yet, if not used responsibly, Big data can increase economic inequality and affirm systemic bias, polarize rather than democratize, and deny opportunities rather than improve access. Worse yet, all this can be done in a way that is non-transparent and defies public scrutiny.

Big data impacts individuals, groups and society as a whole. Because of the central role played by this technology, it must be used *responsibly* – in accordance with the ethical and moral norms that govern our society, and adhering to the appropriate legal and policy frameworks. And as journalists [3], legal and policy scholars [1, 2] and governments [4, 5] are calling for algorithmic fairness and greater insight into data-driven algorithmic processes, there is an urgent need to define a broad and coordinated computer science research agenda in this area. The primary goal of the Dagstuhl Seminar "Data, Responsibly" was to make progress towards such an agenda.

The seminar brought together academic and industry researchers from several areas of computer science, including a broad representation of data management, but also data mining, security/privacy, and computer networks, as well as social sciences researchers, data journalists, and those active in government think-tanks and policy initiatives. The problem we aim to address is inherently transdisciplinary. For this reason, it was important to have input from policy and legal scholars, and to have representation from multiple areas within computer science. We were able to attract a mix of European, North American, and South American participants. Out of 39 participants, 10 were women.

Specific goals of the seminar were to:

- assess the state of data analysis in terms of fairness, transparency and diversity;
- identify new research challenges;
- develop an agenda for computer science research in responsible data analysis and use, with a particular focus on potential high-impact contributions from the data management community;
- solicit perspectives on the necessary education efforts, and on responsible research and innovation practices.

The seminar included technical talks and break-out sessions. Technical talks were organized into themes, which included fairness and diversity, transparency and accountability, tracking and transparency, personal information management, education, and responsible research and innovation. Participants suggested topics for seven working groups, which met over one or multiple days.

The organizers felt that the seminar was very successful – ideas were exchanged, discussions were lively and insightful, and we are aware of several collaborations that were started as a result of the seminar. The participants and the organizers all felt that the topic of the seminar is broad, fast moving and extremely important, and that it would be beneficial to hold another seminar on this topic in the near future.

Details about the program are contained in the remainder of this document.

References

1 Kate Crawford. Artificial Intelligence's White Guy Problem. The New York Times, June 25, 2016.
2 Kate Crawford and Ryan Calo. There is a blind spot in AI research. Nature / Comment 538(7625), October 13, 2016.
3 Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias. ProPublica, May 23, 2016.
4 Executive Office of the President, The White House. Big Data: Seizing Opportunities, Preserving Values. May 2014
5 Parliament and Council of the European Union. General Data Protection Regulation. 2016

2 Table of Contents

Working groups

3 Motivation and overview

3.1 Data, Responsibly: An Overview

Julia Stoyanovich (Drexel University – Philadelphia, US), Serge Abiteboul (ENS – Cachan, FR), and Gerome Miklau (University of Massachusetts – Amherst, US)

Main reference J. Stoyanovich, S. Abiteboul, G. Miklau, "Data Responsibly: Fairness, Neutrality and Transparency in Data Analysis (Tutorial)," in Proc. of the 19th Int'l Conf. on Extending Database Technology (EDBT'16), pp. 718–719, OpenProceedings.org, 2016.
URL http://dx.doi.org/10.5441/002/edbt.2016.103

The first talk of this seminar was a tutorial that surveyed dimensions of responsible data analysis and use, and set the stage for the technical talks and discussions. This presentation was based in part on a recent EDBT tutorial [1].

References

1 J. Stoyanovich, S. Abiteboul, G. Miklau: *D*ata Responsibly: Fairness, Neutrality and Transparency in Data Analysis, Tutorial, EDBT 2016.

3.2 Big Data's Disparate Impact

Solon Barocas (Microsoft – New York, US)

Joint work of Barocas, Solon; Selbst, Andrew
Main reference S. Barocas, A. Selbst, "Big Data's Disparate Impact", California Law Review, Vol. 104, no. 3 (June 2016), pp. 671–732, 2016.
URL http://dx.doi.org/10.15779/Z38BG31

Advocates of algorithmic techniques like data mining argue that these techniques eliminate human biases from the decision-making process. But an algorithm is only as good as the data it works with. Data is frequently imperfect in ways that allow these algorithms to inherit the prejudices of prior decision makers. In other cases, data may simply reflect the widespread biases that persist in society at large. In still others, data mining can discover surprisingly useful regularities that are really just preexisting patterns of exclusion and inequality. Unthinking reliance on data mining can deny historically disadvantaged and vulnerable groups full participation in society. Worse still, because the resulting discrimination is almost always an unintentional emergent property of the algorithm's use rather than a conscious choice by its programmers, it can be unusually hard to identify the source of the problem or to explain it to a court.

This talk examines these concerns through the lens of American antidiscrimination law – more particularly, through Title VII's prohibition of discrimination in employment. In the absence of a demonstrable intent to discriminate, the best doctrinal hope for data mining's victims would seem to lie in disparate impact doctrine. Case law and the Equal Employment Opportunity Commission's Uniform Guidelines, though, hold that a practice can be justified as a business necessity when its outcomes are predictive of future employment outcomes, and data mining is specifically designed to find such statistical correlations. Unless there is a reasonably practical way to demonstrate that these discoveries are spurious, Title VII would appear to bless its use, even though the correlations it discovers will often reflect historic patterns of prejudice, others' discrimination against members of protected groups, or flaws in the underlying data

Addressing the sources of this unintentional discrimination and remedying the corresponding deficiencies in the law will be difficult technically, difficult legally, and difficult politically. There are a number of practical limits to what can be accomplished computationally. For example, when discrimination occurs because the data being mined is itself a result of past intentional discrimination, there is frequently no obvious method to adjust historical data to rid it of this taint. Corrective measures that alter the results of the data mining after it is complete would tread on legally and politically disputed terrain. These challenges for reform throw into stark relief the tension between the two major theories underlying antidiscrimination law: anticlassification and antisubordination. Finding a solution to big data's disparate impact will require more than best efforts to stamp out prejudice and bias; it will require a wholesale reexamination of the meanings of "discrimination" and "fairness".

4 Fairness and diversity

4.1 Fairness Through Awareness and Learning Fair Representations: A Tutorial

Michael Hay (Colgate University – Hamilton, US)

This talk is a tutorial on two recent works on the problem of fairness in classification. The first work, Fairness Through Awareness [1], offers a framework for fair classification that is based on the principle that individuals who are similar for the purpose of the classification task should be treated similarly, and presents an algorithm for maximizing utility subject to the fairness constraint. The second work, Learning Fair Representations [2], formulates fairness as an optimization problem of finding a representation of the data that encodes the data as well as possible while obfuscating information that must be obscured to achieve fairness. The tutorial highlights themes such as individual vs. group fairness, fairness through awareness vs. obfuscation, and formal frameworks vs. empirical assessments.

References

1 Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel Fairness through awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, pp. 214–226. ACM, 2012, http://dx.doi.org/10.1145/2090236.2090255
2 Richard Zemel, Yu Wu,Kevin Swersky, Toniann Pitassi, and Cynthia Dwork Learning Fair Representations. In Proceedings of the International Conference on Machine Learning, pp. 325–333. 2013

4.2 An Axiomatic Framework for Fairness

Suresh Venkatasubramanian (University of Utah – Salt Lake City, US), Carlos Scheidegger, and Sorelle Friedler (Haverford College, US)

Joint work of S. A. Friedler, C. Scheidegger, S. Venkatasubramanian
Main reference S. A. Friedler, C. Scheidegger, S. Venkatasubramanian, "On the (im)possibility of fairness," arXiv:1609.07236v1 [cs.CY], 2016.
URL https://arxiv.org/abs/1609.07236v1

What does it mean for an algorithm to be fair? Different papers use different notions of algorithmic fairness, and although these appear internally consistent, they also seem mutually incompatible. We present a mathematical setting in which the distinctions in previous papers can be made formal. In addition to characterizing the spaces of inputs (the "observed" space) and outputs (the "decision" space), we introduce the notion of a construct space: a space that captures unobservable, but meaningful variables for the prediction.

We show that in order to prove desirable properties of the entire decision-making process, different mechanisms for fairness require different assumptions about the nature of the mapping from construct space to decision space. The results in this work imply that future treatments of algorithmic fairness should more explicitly state assumptions about the relationship between constructs and observations.

4.3 What is Fairness Anyway? Interdisciplinary Concepts and Data Science

Bettina Berendt (KU Leuven, BE)

Main reference B. Berendt, S. Preibusch, "Better decision support through exploratory discrimination-aware data mining: foundations and empirical evidence," Artif. Intell. Law, 22(2): 175–209, 2014.
URL http://dx.doi.org/10.1007/s10506-013-9152-0
Main reference B. Berendt, M. Büchler, G. Rockwell, "Is it Research or is it Spying? Thinking-Through Ethics in Big Data AI and Other Knowledge Sciences," KI, 29(2):223–232, 2015.
URL http://dx.doi.org/10.1007/s13218-015-0355-2

It is by now a truism that data mining algorithms "discriminate". Whether we consider any particular criterion or effect to be a differentiation or an undesirable (e.g. unlawful) "discrimination in the narrow sense", is a second question. Similarly, one needs to investigate the non-identical notions of "non-discrimination" and "fairness".

Thus, conceptual issues arise even at the start of any process of being more responsible about data, and addressing them requires an interdisciplinary approach. Still, no data mining algorithm by itself is discriminatory – it can only become so when deployed in a context. This talk builds on our work on investigating discrimination-aware data mining (DADM) in contexts that involve human decision makers. I give a brief overview of our proposal of an interactive, exploratory DADM and an empirical study we did of such decisions. I then present five challenges that cannot be tackled by today's formalisms for DADM or "fairness-aware data mining": vicious cycles of one form of discrimination leading into another, the question of how to translate the Aristotelian principle of equality into a data framework, intersectionality and the emergence of new concepts, the perpetuation of pernicious concepts and how it gets baked into data-based decision making, and the search for causes. I conclude with an outlook on next-generation tools and research approaches.

The slides of the talk are available at [1].

References

1 Bettina Berendt. *What is fairness anyway? Interdisciplinary concepts and data science*. Presentation at the Dagstuhl Seminar "Data, responsibly". Dagstuhl, 18 July 2016. https://people.cs.kuleuven.be/~bettina.berendt/Talks/berendt_2016_07_18.pptx

4.4 Segregation Discovery

Salvatore Ruggieri (University of Pisa, IT)

Joint work of Baroni, Alessandro
Main reference A. Baroni, S. Ruggieri, "Segregation Discovery in a Social Network of Companies," in Proc. of the 14th Int'l Symp. on Intelligent Data Analysis (IDA'15), LNCS, Vol. 9385, pp. 37–48, Springer, 2015.
URL http://dx.doi.org/10.1007/978-3-319-24465-5_4

The term segregation refers to restrictions on the access of people to each other. People are partitioned into two or more groups on the grounds of personal or cultural traits that can foster discrimination, such as gender, age, ethnicity, income, skin color, language, religion, political opinion, membership of a national minority, etc. Contact, communication, or interaction among groups are limited by their physical, working or socio-economic distance.

We introduce a framework for a data-driven analysis of segregation of minority groups in social networks, and challenge it on a complex scenario. The framework builds on quantitative measures of segregation, called segregation indexes, proposed in the social science literature. The segregation discovery problem consists of searching sub-graphs and sub-groups for which a reference segregation index is above a minimum threshold. A search algorithm is devised that solves the segregation problem based on frequent itemset mining. The framework is challenged on the analysis of segregation of social groups in the boards of directors of the real and large network of Italian companies connected through shared directors.

Relationships among segregation, discrimination, and diversity are also highlighted.

4.5 Diversity: Why, What, How

Evaggelia Pitoura (University of Ioannina, GR) and Marina Drosou (Hellenic Police – Athens, GR)

In this talk, we first present a brief overview of data diversification. We discuss different diversification interpretations, namely, based on coverage, content dissimilarity and novelty, as well as, various algorithmic approaches. Then, we present our work on r-DisC diversification. r-DisC diversification locates diverse subsets of results in a way such that each item in the result is represented by a similar item in the diverse subset and the items in the diverse subset are dissimilar to each other. We also show various extensions of our basic model. Finally, we discuss some issues in social networks and opinion diversity, such as, homophily, opinion formation and fairness.

5 Transparency and accountability

5.1 Accountable Algorithms

Solon Barocas (Microsoft – New York, US)

Joint work of Joshua A. Kroll, Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, Harlan Yu
Main reference J. A. Kroll, J. Huey, S. Barocas, E. W. Felten, J. R. Reidenberg, D. G. Robinson, H. Yu, "Accountable Algorithms", University of Pennsylvania Law Review, forthcoming; pre-print available at SSRN.
URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2765268

Many important decisions historically made by people are now made by computers. Algorithms count votes, approve loan and credit card applications, target citizens or neighborhoods for police scrutiny, select taxpayers for an IRS audit, and grant or deny immigration visas.

The accountability mechanisms and legal standards that govern such decision processes have not kept pace with technology. The tools currently available to policymakers, legislators, and courts were developed to oversee human decision-makers and often fail when applied to computers instead: for example, how do you judge the intent of a piece of software? Additional approaches are needed to make automated decision systems – with their potentially incorrect, unjustified or unfair results – accountable and governable. This Article reveals a new technological toolkit to verify that automated decisions comply with key standards of legal fairness.

We challenge the dominant position in the legal literature that transparency will solve these problems. Disclosure of source code is often neither necessary (because of alternative techniques from computer science) nor sufficient (because of the complexity of code) to demonstrate the fairness of a process. Furthermore, transparency may be undesirable, such as when it permits tax cheats or terrorists to game the systems determining audits or security screening.

The central issue is how to assure the interests of citizens, and society as a whole, in making these processes more accountable. This Article argues that technology is creating new opportunities – more subtle and flexible than total transparency – to design decision-making algorithms so that they better align with legal and policy objectives. Doing so will improve not only the current governance of algorithms, but also – in certain cases – the governance of decision-making in general. The implicit (or explicit) biases of human decision-makers can be difficult to find and root out, but we can peer into the "brain" of an algorithm: computational processes and purpose specifications can be declared prior to use and verified afterwards.

The technological tools introduced in this Article apply widely. They can be used in designing decision-making processes from both the private and public sectors, and they can be tailored to verify different characteristics as desired by decision-makers, regulators, or the public. By forcing a more careful consideration of the effects of decision rules, they also engender policy discussions and closer looks at legal standards. As such, these tools have far-reaching implications throughout law and society.

Part I of this Article provides an accessible and concise introduction to foundational computer science concepts that can be used to verify and demonstrate compliance with key standards of legal fairness for automated decisions without revealing key attributes of the decision or the process by which the decision was reached. Part II then describes how these techniques can assure that decisions are made with the key governance attribute of procedural regularity, meaning that decisions are made under an announced set of rules consistently applied in each case. We demonstrate how this approach could be used to redesign and resolve issues with the State Department's diversity visa lottery. In Part III, we go further and explore how other computational techniques can assure that automated decisions preserve fidelity to substantive legal and policy choices. We show how these tools may be used to assure that certain kinds of unjust discrimination are avoided and that automated decision processes behave in ways that comport with the social or legal standards that govern the decision. We also show how algorithmic decision-making may even complicate existing doctrines of disparate treatment and disparate impact, and we discuss some recent computer science work on detecting and removing discrimination in algorithms, especially in the context of big data and machine learning. And lastly in Part IV, we propose an agenda to further synergistic collaboration between computer science, law and policy to advance the design of automated decision processes for accountability.

5.2 Algorithmic Accountability and Transparency

Nicholas Diakopoulos (University of Maryland – College Park, US)

Main reference N. Diakopoulos, "Accountability in Algorithmic Decision Making," Communications of the ACM, 59(2):56–62, ACM, 2016.
URL http://dx.doi.org/10.1145/2844110

As journalism shifts into the 21st century, the opportunities for reinventing the ways that news stories are found and told using computing and data are practically endless. But perhaps even more substantial are the new ways in which computation and algorithms are coming to adjudicate decisions in nearly all facets of industry and government. Algorithmic accountability reporting is a new form of computational journalism that is emerging to apply the core journalistic functions of watchdogging and investigative reporting to algorithms. In this talk I will discuss how algorithmic accountability reporting is used by journalists as a method for articulating the power structures, biases, and influences that computational artifacts play in society. I will trace various legal, technical, and regulatory challenges that remain, offering new openings for the development of tools. Finally, I will discuss the mandate for transparency of algorithms and proffer for discussion an initial transparency standard that delineates the dimensions of algorithms in use by industry or government that might be disclosed.

5.3 Auditing Black-box Models

Sorelle Friedler (Haverford College, US)

Joint work of Philip Adler, Casey Falk, Sorelle A. Friedler, Gabriel Rybeck, Carlos Scheidegger, Brandon Smith, Suresh Venkatasubramanian
Main reference P. Adler, C. Falk, S. A. Friedler, G. Rybeck, C. Scheidegger, B. Smith, S. Venkatasubramanian, "Auditing Black-box Models by Obscuring Features," arXiv:1602.07043v1 [stat.ML] , 2016.
URL http://arxiv.org/abs/1602.07043v1

Data-trained predictive models see widespread use, but for the most part they are used as black boxes which output a prediction or score. It is therefore hard to acquire a deeper understanding of model behavior, and in particular how different features influence the model prediction. This is important when interpreting the behavior of complex models, or asserting that certain problematic attributes (like race or gender) are not unduly influencing decisions.

In this talk, I present a technique for auditing black-box models, which lets us study the extent to which existing models take advantage of particular features in the dataset, without knowing how the models work. Our work focuses on the problem of indirect influence: how some features might indirectly influence outcomes via other, related features. As a result, we can find attribute influences even in cases where, upon further direct examination of the model, the attribute is not referred to by the model at all.

Our approach does not require the black-box model to be retrained. This is important if (for example) the model is only accessible via an API, and contrasts our work with other methods that investigate feature influence like feature selection. We present experimental evidence for the effectiveness of our procedure using a variety of publicly available datasets and models. Not presented, we also validate our procedure using techniques from interpretable learning and feature selection, as well as against other black-box auditing procedures.

5.4 Revealing Algorithmic Rankers

Gerome Miklau (University of Massachusetts – Amherst, US) and Julia Stoyanovich (Drexel University – Philadelphia, US)

Joint work of Julia Stoyanovich, Gerome Miklau, Ellen P. Goodman
Main reference J. Stoyanovich, E. P. Goodman, "Revealing Algorithmic Rankers," Freedom to Tinker, Nov. 2016.
URL https://freedom-to-tinker.com/blog/jstoyanovich/revealing-algorithmic-rankers/

ProPublica's story on "machine bias" in an algorithm used for sentencing defendants amplified calls to make algorithms more transparent and accountable. It has never been more clear that algorithms are political and embody contested choices, and that these choices are largely obscured from public scrutiny. We see it in controversies over Facebook's newsfeed, or Google search results, or Twitter's trending topics. Policymakers are considering how to operationalize "algorithmic ethics" and scholars are calling for accountable algorithms.

One kind of algorithm that is at once especially obscure, powerful, and common is the ranking algorithm. Algorithms rank individuals to determine credit worthiness, desirability for college admissions and employment, and compatibility as dating partners. They encode ideas of what counts as the best schools, neighborhoods, and technologies. Despite their importance, we actually can know very little about why this person was ranked higher than another in a dating app, or why this school has a better rank than that one. This is true even if we have access to the ranking algorithm, for example, if we have complete knowledge about the factors used by the ranker and their relative weights, as is the case for US News ranking of colleges. In this blog post, we argue that syntactic transparency, wherein the rules of operation of an algorithm are more or less apparent, or even fully disclosed, still leaves stakeholders in the dark: those who are ranked, those who use the rankings, and the public whose world the rankings may shape.

In this talk we discuss the reasons for opacity in ranking algorithms, and the corresponding harms that this opacity brings. We give examples of these issues in using rankings of US colleges and academic departments. We go on to outline directions for future work that would make rankings interpretable.

5.5 Computational Fact Checking

Cong Yu (Google – New York, US)

Joint work of Bill Adair, Chengkai Li, Jun Yang, Cong Yu

This talk describes the process through which reporters perform fact checking on statements made by politicians and the various challenges facing the reporters. The three main stages in the process are finding the claims, checking the claims, and distributing the reviews. For each stage, I illustrate some recent works that aim at computationally assisting the reporters, as well as more technical challenges to be addressed for the ultimate holy grail of automatic fact checking. This talk is also a call-for-action for database and algorithm researchers to work in this socially important area.

6 Tracking and transparency

6.1 Online Tracking and Transparency

Claude Castelluccia (INRIA – Grenoble, FR)

URL https://myrealonlinechoices.inrialpes.fr/

In the last few years, as a result of the proliferation of intrusive and privacy-invading ads, the use of ad-blockers and anti-tracking tools have become widespread. As of the second quarter of this year, 16% of online Americans, about 45 million people, had installed ad-blocking software, according to PageFair 2015 report. Meanwhile, 77 millions Europeans are blocking ads. All this accounts globally for $21.8 billion worth of blocked ads. The Internet economy is in danger since ads fuel the free content and services over Internet.

We believe that Adblockers are only a short-term solution, and that better tools are necessary to solve this problem in the long term. Most users are not against ads and are actually willing to accept some ads to help websites. However, users want more control and transparency about the ads that they want to receive, and about the way they are tracked and profiled on the Internet.

As opposed to existing ad blockers that take a binary approach (i.e., block everything if you install them or block nothing otherwise), MyRealOnlineChoices project aims to provide users with the right tools that allow users to make fine-grained choices about their privacy and the ads that they want to receive. Our tools allow users to choose on which sites (more specifically, on which categories of sites) they want to block the trackers. For example, a user can choose to block the trackers on sites related to health or religion, but may choose not to block the trackers on sites related to sports or news. Similarly, a user might want to block ads that are targeted on some categories that he considers sensitive. Our tools provide this type of control to the users.

As trust between users and online entities is the key here, our project starts with transparency as the first key feature. We need tools that provide more transparency to users by indicating if an ad is retargeted or is delivered to users based on their interests or not. The ultimate goal is to enforce the user choices while sustaining the ad economy of the Internet. And thanks to this transparency feature, users can be aware of what is going on with their browsing data, and therefore can make an informed decision.

6.2 Tracing Information Flows Between Ad Exchanges Using Retargeted Ads

Christo Wilson

Joint work of Muhammad Ahmad Bashir, Sajjad Arshad, William Robertson, Christo Wilson
Main reference M. A. Bashir, S. Arshad, W. Robertson, C. WIlson, "Tracing Information Flows Between Ad Exchanges Using Retargeted Ads," in Proc. of the 25th USENIX Security Symp., pp. 481–496, USENIX Association, 2016.
URL https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/bashir

Numerous surveys have shown that Web users are seriously concerned about the loss of privacy associated with online tracking. Alarmingly, these surveys also reveal that people are also unaware of the amount of data sharing that occurs between ad exchanges, and thus underestimate the privacy risks associated with online tracking.

In reality, the modern ad ecosystem is fueled by a flow of user data between trackers and ad exchanges. Although recent work has shown that ad exchanges routinely perform cookie matching with other exchanges, these studies are based on brittle heuristics that cannot detect all forms of information sharing, especially under adversarial conditions.

In this study, we develop a methodology that is able to detect client- and server-side flows of information between arbitrary ad exchanges. Our key insight is to leverage retargeted ads as a mechanism for identifying information flows. Intuitively, our methodology works because it relies on the semantics of how exchanges serve ads, rather than focusing on specific cookie matching mechanisms. Using crawled data on 35,448 ad impressions, we show that our methodology can successfully categorize four different kinds of information sharing between ad exchanges, including cases were existing heuristic methods fail.

References

1 Muhammad Ahmad Bashir and Sajjad Arshad and William Robertson and Christo WIlson. *Tracing Information Flows Between Ad Exchanges Using Retargeted Ads*. In Proceedings of Usenix Security, Austin TX, August 2016.

6.3 Quantifying Search Engine Bias

Krishna P. Gummadi (MPI-SWS – Saarbrücken, DE)

Joint work of Juhi Kulshrestha, Motahhareh Eslami, Saptarshi Ghosh, Krishna P. Gummadi, Karrie Karahalios
Main reference J. Kulshrestha, M. Eslami, J. Messias, M.B. Zafar, S. Ghosh, K.P. Gummadi, K. Karahalios, "Quantifying Search Bias: Investigating Sources of Bias for Political Searches in Social Media," to appear in Proc. of the 20th ACM Conf. on Computer-Supported Cooperative Work and Social Computing(CSCW'17); pre-print available from author's webpage.
URL www.mpi-sws.org/ juhi/search-bias-cscw-2017-pre-print.pdf

Search systems in online social media sites are frequently used to find information about ongoing events and people. For topics with multiple competing perspectives, such as political events or political candidates, bias in the top ranked results significantly shapes public opinion. However, bias does not emerge from an algorithm alone. It is important to distinguish between the bias that arises from the data that serves as the input to the ranking algorithm and the bias that arises from the ranking algorithm itself. In this talk, I will propose a framework to quantify these distinct biases and apply this framework to politics-related queries on Twitter. We found that both the input data and the ranking algorithm contribute significantly to produce varying amounts of bias in the search results and in different ways. I will discuss the consequences of these biases and propose mechanisms to signal this bias in social search systems interfaces.

6.4 Seeing through Website Privacy Policies

Rishiraj Saha Roy (MPI für Informatik – Saarbrücken, DE)

A number of online privacy concerns involving data sharing and unexpected ad recommendations can be mitigated if users understand website privacy policies better. Unfortunately these policies are usually very long since they contain legal documentation, and have to cover several corner cases. The short message on cookies that we get when we visit a new website inside Europe is also generally not enough. Wilson et al. (2016) state that the three key things that a user is really concerned about in a website's privacy policy are whether it collects personal information, shares that personal information, and how easy it is to delete this personal information. The authors find that the average Web user is indeed able to find answers to most of these concerns inside the policies, and it is possible to automatically extract relevant excerpts from the long privacy policies using regular expressions and weighted term matching. With this knowledge, we can trivially show an extractive summarization of the privacy policy focused on the relevant aspects to the user instead of the short cookie message, and the list of trackers if any using tools like Ghostery. But going further, we can get the preferred privacy policy from the user as defined by answers to the questions earlier, and then try to automatically reason and find the likely answers. Using this knowledge, we can then check the compliance of the site's policy to the user's preference, and then suggest alternatives. An unobtrusive way of suggesting alternatives would be to show policy-compliant websites on the search results' page itself, as users usually end up on these sites through a search engine. So, in a post-processing step, our privacy advisor can fetch the privacy policies of the top-k sites (k = 5 or 10, say), and check the compliance statuses before presenting to the user, and the user can accordingly choose his/her preferred site(s). This specific use case is applicable of hundreds of usual websites that we visit for day-to-day information like flight fare and insurance comparisons, but not really to big players like Google, Facebook or Amazon. There are quite a few research challenges here: sometimes the clarification on a privacy aspect is unclear even to humans, and sometimes there are disagreements between average Web users and legal experts. The automated reasoning about policy compliance using NLP techniques is non-trivial, and finally not every visit to these websites is through a search engine, so we have to figure out how best to show relevant alternatives in such cases.

6.5 Collect it All: Why Bulk Surveillance Works

Nicholas Weaver (ICSI – Berkeley, US)

Joint work of NSA

A big driving force behind the “collect it all” mentality is that bulk surveillance techniques, both for the private sector and intelligence services, both works and is economically feasible. When developing a surveillance system, no matter the purpose, there is usually a requirement to have the capability to target anyone. As a corollary there is also a huge desire for retrospective capabilities, since the data collector doesn’t know until tomorrow what he wished to save today. This leads to architectures which effectively collect data on everybody and then, subsequently, select for the actual information of interest.

Such architectures are also remarkably affordable and, thanks to modern big data techniques, scale linearly. Ranging in size from just a couple of racks to a large facility, its straightforward to match both the data budget and computational budget needed to collect, retain, and search information on everybody. $100M in hardware could maintain a 10MB dossier on everybody on the planet, It may be the case that "collect it all" is simply unstoppable because it is both effective and affordable.

6.6 Tracking, Targeting, Rating, Discriminating based on Social Media: Risk Measures for User Guidance

Gerhard Weikum (MPI für Informatik – Saarbrücken, DE)

Joint work of Joanna Biega, Krishna Gummadi, Gerhard Weikum

In this talk I introduce the R-Susceptibility model which captures sensitive topics in online communities and the exposure or risk of individual users incurred by their posts. Topics are captured by latent embeddings, and sensitive topics are identified by crowdsourcing. A user’s risk is quantified by distance measures between the user’s post or entire posting history and the topic of interest. The R-Susceptibility model allows ranking users with regard to an adversary that targets the top-k users on some sensitive topic. Based on these methods, we envision a personalized tool that can alert users, explain risks and possible countermeasures, and guides users.

7 Personal information management

7.1 Managing your Personal Information

Serge Abiteboul (ENS – Cachan, FR) and Amélie Marian (Rutgers University – Piscataway, US)

Personal information is constantly produced and stored by a large number of sources and services. While in the past, users would store all their data in physical form or on their local machines, the advent of the cloud has made this impossible. Personal information is fragmented in multiple heterogeneous systems, making it difficult to access, control, and exploit for personal use. In this talk, we make the case for the need for Personal Information Management Systems (PIMS), (cloud-based) systems that manages all the information of a person. Recent technological advances and societal pressures are enabling new interesting applications for PIMS. We discuss a subset of these applications: data integration, personal search, and personal knowledge management, and their potential to improve users' lives by giving them back control over their own information.

7.2 Small Data Metadata

Arnaud Sahuguet (Cornell Tech NYC, US)

Main reference Arnaud Sahuguet, "Small Data Metadata," 2016.
URL https://medium.com/@sahuguet/small-data-metadata-1ff922fa6d14#.rmzh2c9kh

Small data are the digital traces that individuals generate as a byproduct of daily activities, such as sending e-mail or exercising with fitness trackers.

We advocate for the need to standardize metadata about small data and build tools to create, manage, process and reason about it. Not only is small data metadata critical to foster the small data ecosystem of consumers and producers, it is also essential to offer some necessary guarantees – like privacy – inherent to the data itself.

7.3 Empowering Personal Data Management using Secure Hardware

Benjamin Nguyen (INSA – Bourges, FR)

Joint work of Nicolas Anciaux, Philippe Bonnet, Luc Bouganim, Benjamin Nguyen, Philippe Pucheral, Iulian Sandu-Popa
Main reference N. Anciaux, P. Bonnet, L. Bouganim, B. Nguyen, I. S. Popa, P. Pucheral, "Trusted Cells: A Sea Change for Personal Data Services," in Proc. of the 6th Biennial Conf. on Innovative Data Systems Research (CIDR'13), 2013.
URL http://cidrdb.org/cidr2013/Papers/CIDR13_Paper68.pdf

How do you keep a secret about your personal life in an age where your daughter's glasses record and share everything she senses, your wallet records and shares your financial transactions, and your set-top box records and shares your family's energy consumption? Your personal data has become a prime asset for many companies around the Internet, but can you avoid – or even detect – abusive usage?

Today, there is a wide consensus that individuals should have increased control on how their personal data is collected, managed and shared. Yet there is no appropriate technical solution to implement such personal data services: centralized solutions sacrifice security for innovative applications, while decentralized solutions sacrifice innovative applications for security. In this presentation, we argue that the advent of secure hardware in all personal IT devices, at the edges of the Internet, could trigger a sea change.

We introduce PlugDB, a personal data server running on a secure portable tokens which forms a global, decentralized data platform that provides security yet enables innovative applications. We describe this platform, called asymmetric architecture, because it is composed on the one hand of a large number of low power, low availability, high trust devices, and on the other hand high power, 24/7 but low to no trust cloud type infrastructure. Finally, we define a range of challenges for future research.

References

1 Cuong-Quoc To, Benjamin Nguyen, Philippe Pucheral. *Private and Scalable Execution of SQL Aggregates on a Secure Decentralized Architecture*, TODS, to appear, 2016.
2 Nicolas Anciaux, Saliha Lallali, Iulian Sandu Popa, Philippe Pucheral. *A Scalable Search Engine for Mass Storage Smart Objects*. VLDB 2015

8 Education, responsible research and innovation

8.1 Science Data, Responsibly

Bill Howe (University of Washington – Seattle, US)

Joint work of Maxim Gretchkin, Hoifung Poon, Poshen Lee

There is a reproducibility crisis in science: the number of retractions are increasing year to year, public trust is low, and a number of reproducibility studies across fields have shown dismal results. The incentive structures in science have increased pressure to achieve results at all costs, and new technology has made it easier to substitute exploratory research for controlled experiments.

We see two complementary solutions: In education, we advocate incorporating a rigorous ethics program into data science curricula, emphasizing case studies that do not readily admit technical solutions. In technology, we advocate systems research to enforce statistical checks, avoid multiple hypothesis testing issues, and ensure curation of public datasets.

As an example, we describe a project in computational curation that provides a first step toward automatic verification of scientific claims. Using a public repository of microarray data, we show that co-trained models on the human-provided metadata and the content of the dataset itself can significantly improve the quality of the labels over the state of the art methods, making thousands of new datasets available for reproducibility studies.

8.2 Research and Education in Data Science and Responsible Use: Challenges and Opportunities

Chaitanya Baru (NSF – Arlington, US)

This talk will provide an overview of activities supported by the US National Science Foundation (NSF), and recent initiatives in the US Federal Government, that address the issue of responsible use of data. The talk is intended to initiate a discussion on the role that funding agencies could play in supporting a research and education agenda in this area.

8.3 Sustainability Research: Promoting Transparency and Accountability for Decision Makers

Claudia Bauzer Medeiros (UNICAMP – Campinas, BR)

Sustainability is a transdisciplinary research domain that is strongly dependent on the analysis of big data, at multiple space and time scales. In a broad sense, it can be seen as an effort to improve people's lives without compromising the planet's limited resources, from a micro point of view (a single person) to a macro perspective (the Earth, and the biosphere).

Sustainability studies cover a vast range of subjects, such as health, pollution and climate change, biodiversity, inequality or education. As a consequence, data handled are widely heterogeneous, e.g., concerning records about an individual, or environmental measurements, or observations of species.

The talk will discuss a few of the research challenges for big data analysis in sustainability via a real use case in Brazil, and the intrinsic scientific, economic, social and political issues. It will emphasize the aspects of transparency, auditability and accountability for decision making in this context.

8.4 Practising Responsible Data Practices through Data Ethics Education

H. V. Jagadish (University of Michigan – Ann Arbor, US)

URL https://www.edx.org/course/data-science-ethics-michiganx-ds101x

We will get responsible data practices only if data practitioners are responsible, and data practitioners will be responsible only if they know how. For these reasons, it is critical to (1) Make Data Scientists aware of ethical issues regarding data so that they at least try to do the right thing, AND (2) Empower them with the tools to do the right thing with minimum burden.

Imperative (1) means that we need Data Ethics training as an integral part of Data Science training. One suggested starting point is a recent MOOC on EdX. Imperative (2) means that we need research to develop algorithms and tools that can effectively implement policies that we societally decide are the ones we would like to adopt.

There is an additional question of what policies we should adopt. This requires social consensus, at least across certain segments of society. An informed consensus can only be reached with good education. So the need for Data Ethics education actually extends beyond just Data Science practitioners and to society at large.

Finally, we note that there are many challenging issues at the boundaries – defining what exactly is OK, sociological issues in developing consensus, political concerns regarding laws and regulations, and so on. However, there is a great deal that we can all agree about today – stuff that is not at the boundaries. There is urgent pressure to get this to practice.

8.5 Values, Algorithm Design, and Collaboration

Kristene Unsworth (Drexel University – Philadelphia, US)

Algorithms and the results they provide appear to many, outside our fields to be objective. We know this is not the case and that the values of data scientists and anyone working with algorithms are reflected in these designs. Because of this, it is important to acknowledge the importance of ethics in the work we do and in society Ethics are about action and our values drive us. This presentation discussed ongoing research into the role of values in algorithm design and within teams. The work / algorithm-related values of the workshop participants were also discussed and highlighted relation to early research findings. These included:

- technological progress leads to social progress
- correcting information asymmetries
- have an impact on the community
- knowledge
- clarity / insight
- curiosity
- sustainable business model
- intellectual outrage
- technological solutions
- awareness of background context

8.6 Privacy, Transparency and Education

Gerald Friedland (ICSI – Berkeley, US)

Joint work of Gerald Friedland, Dan Garcia, Julia Bernd, Serge Egelmann, Jaeyoung Choi
URL http://www.teachingprivacy.org

Decisions about data sharing begin early in somebody's life. For example, when one uploads an image to the Internet, the decision might be whether to post the image or not given the other people that can be identified in that photograph. In our times, decisions like that start to arise in teenage years. On the other hand, current curricula do not cater to the new responsibilities, not even on the level of University education for engineers. In my talk at Dagstuhl, I presented some of the new issues that arise, including cybercasing [1], data exploitation, and privacy concern followed by a presentation of the teaching resources http://www.teachingprivacy.org.

The Teaching Privacy project is an NSF-sponsored collaboration between the International Computer Science Institute and the University of California-Berkeley. The project aims to empower high school students and college undergrads in making informed choices about privacy, by building a set of educational tools and hands-on exercises to help teachers demonstrate what happens to personal information on the Internet, and what the effects of sharing information can be.

Current computer-science curricula at high schools and colleges usually include an abundance of material on data-retrieval methods and how to improve them, but rarely make room for discussion of the potential negative impact of these technologies. Among the groups most affected by those negative impacts are high-school students; they are the most frequent users of social-networking sites and apps, but often do not have a full understanding of the potential consequences their current online activities might have later in their lives. For example, a Facebook posting that a high-scooters friends think is cool might be seen by a much larger audience than she expected – including perhaps future employers who would not think it was so cool. In addition, not understanding – or not thinking about – the consequences of posting often leads to over-sharing information about other people, including friends and relatives.

The Teaching Privacy is organized around 10 basic principles. Each principle is underpinned with technical explanations, anecdotes, apps, videos, news links, exercises, discussion items, and a guideline for teachers. The teacher's guidelines (TROPE – Teacher's Resources for Online Privacy Education) follow the paradigm of the 5 E's [3]. Learning objectives are outlined clearly and the whole site is licensed under Creative Commons 0, allowing a teacher to cherry pick from the website and creating their own curriculum [2].

This work was supported by funding provided to the International Computer Science Institute by the National Science Foundation, through grants CNS-1065240 and DGE-1419319, and by the Broadband Technology Opportunities Program through the California Connects program. Additional support comes from funding provided to the University of California Berkeley through NSF grants EEC-1405547 and CCF-0424422 and through the IISME Summer Fellowship Program.

References

1 G. Friedland, R. Sommer: *C*ybercasing the Joint: On the Privacy Implications of Geotagging, Usenix HotSec 2010 at the Usenix Security Conference, Washington DC, August 2010.

2 Julia Bernd, Blanca Gordo, Jaeyoung Choi, Bryan Morgan, Nicholas Henderson, Serge Egelman, Daniel D Garcia, Gerald Friedland: *T*eaching Privacy: Multimedia Making a Difference, IEEE MultiMedia, Vol. 22, Issue 1, pp. 12–19, January 2015.

3 http://enhancinged.wgbh.org/research/eeeee.html

8.7 Teaching Ethical Issues in Data Mining to Undergraduates

Sorelle Friedler (Haverford College, US)

URL http://ww3.haverford.edu/computerscience/courses/cmsc207/

In this talk, I discuss how guiding principles about teaching ethics in data mining to undergraduates were integrated into the design of 100-level and 200-level courses. The main principle underlying these curricular choices is that fairness and ethical considerations should be integrated throughout the course, not sidelined to a single lecture or module. These fairness and ethical issues are discussed throughout the course as they relate to the understanding of real world data (e.g., data errors and choices) and the communication of these choices and assumptions (e.g., error values and contextual assumptions). Conversations with domain experts help to connect the data to its true context and drive home the importance and ethical nature of the choices made.

8.8 Networked Systems Ethics

Ben Zevenbergen (University of Oxford, GB)

URL http://networkedsystemsethics.net/

The Oxford Internet Institute's Ethics in Networked Systems Research project is developing practical guidelines for computer scientists and engineers to assess the ethical implications and social impact of Internet-based projects. These guidelines aim to underpin a meaningful cross-disciplinary conversation between gatekeepers of ethics standards and researchers about the ethical and social impact of technical Internet research projects. The iterative reflexivity methodology guides stakeholders to identify and minimize risks and other burdens, which must be mitigated to the largest extent possible by adjusting the design of the project before data collection takes place. The aim is thus to improve the ethical considerations of individual projects, but also to streamline the proceedings of ethical discussions in Internet research generally. The primary audience for these guidelines are technical researchers (e.g. computer science, network engineering, as well as social science) and gatekeepers of ethics standards at institutions, academic journals, conferences, and funding agencies. It is possible to use these guidelines beyond in academic research in civil society, product development, or otherwise, but these are not the primary audience. Some sections point the reader to other groups – such as the data subjects, lawyers, local peers, etc. – who can also use (parts of) the guidelines to help assess the impact of a project from their expertise or point of view.

9 Lightning talks

9.1 Benchmarking for (Linked) Data Management

Irini Fundulaki (FORTH – Heraklion, GR)

In this talk we discussed the importance of benchmarking and focused on the kinds of benchmarks we need in order to assess the responsibility regarding the use of data in general, by the different service providers: search engines, recommendation services, and social networks among others. Benchmarks will allow us to measure the bias, or neutrality of the aforementioned entities and give people and regulating bodies a good basis on how to act and react regarding the protection and we'll use of their data.

9.2 From Three Laws of Robotics to Five Principles of Big Data?

Wolfgang Nejdl (Leibniz Universität Hannover, DE)

Can we define five principles for Big Data to guide us through future big data applications? If these are general / generic guidelines, yes. If they should constrain big data applications in an effective way, no. The world has become too complex, possible scenarios are two diverse, and different aspects of big data applications are too often conflicting. How can we solve this dilemma? By working on the issues the workshop focused on in an interdisciplinary way, and by not only focusing on efficiency of algorithms, but also on their fairness and related aspects.

9.3 Natural Language Processing, Responsibly

Jannik Strötgen (MPI für Informatik – Saarbrücken, DE)

An important kind of data that should be dealt with responsibly, is unstructured data. There is a huge amount of textual data containing information that is not (yet) available in structured format, and natural language processing techniques can be applied to get structured information out of it. Until few years ago, the main text types that NLP research dealt with were non-personal data such as news corpora, and the main goal has been to enrich the textual content with further information. Nowadays, a lot of NLP research is carried out on personal data such as social media content. Although enriching documents is still a major goal, a lot of research addresses predicting author characteristics. Thus, and because NLP techniques are not just used in research environments, but became mature enough to be applied for all types of (real-time) applications, their output can clearly affect individuals. Besides false extractions and aggregations, further problems occur as language is typically uttered at a specific time and place in a specific context, but extracted and aggregated information extracted from textual data does often not consider this context anymore. Thus, as data in general, unstructured data and NLP techniques should be used and applied responsibly.

10 Working groups

10.1 Structural Bias

Solon Barocas (Microsoft – New York, US), Bettina Berendt (KU Leuven, BE), Michael Hay (Colgate University – Hamilton, US), Amélie Marian (Rutgers University – Piscataway, US), and Gerome Miklau (University of Massachusetts – Amherst, US)

This working group discussed the problem of structural bias in algorithm output. The group agreed that bias was not the best term, but rather that the focus should be on structural discrimination.

The following questions were raised during the discussion: (1) How can we assess structural discrimination? (2) How much of the discrimination comes from the algorithm, how much of it from the input data? (3) If the data is biased, is it the job of the algorithm (operator) to correct structural bias? How? And wouldn't that introduce new bias, that of the algorithm designer/operator? The discussion kept returning to the last question, with various examples, and the group recognized that we were rehashing the traditional issues/questions of affirmative action (fairness w.r.t. to skills or potential, making up for societal bias).

Starting from the model of true space vs. observed space of Friedler et al. [1], an important question is how to identify and mitigate the errors in the input data that may be introducing discrimination. A possibility is to combine multiple observation measurements, and to correct for some of the known bias. One problem with this approach is that many causes of discrimination cannot be identified. In addition, this approach assumes that most of the bias comes from the observation, and not from explicit bias in the data, which is unlikely to be true in practice. Correcting bias in the data is both technically challenging and may have legal ramifications.

References

1 Sorelle A. Friedler, Carlos Scheidegger and Suresh Venkatasubramanian On the (im)possibility of fairness. *arXiv preprint arXiv:1609.07236*, 2016.

10.2 Avoid Reinventing the Wheel

Bettina Berendt (KU Leuven, BE), Solon Barocas (Microsoft – New York, US), Claude Castelluccia (INRIA – Grenoble, FR), Pauli Miettinen (MPI für Informatik – Saarbrücken, DE), Wolfgang Nejdl (Leibniz Universität Hannover, DE), Salvatore Ruggieri (University of Pisa, IT), and Jannik Strötgen (MPI für Informatik – Saarbrücken, DE)

The working group investigated the question of how we can avoid reinventing various wheels when developing more responsible approaches to dealing with data.

We identified different dimensions along which existing wheels should be studied in order to learn from them, in particular: research contents (a topic discussed briefly and repeatedly at the Dagstuhl seminar, albeit outside the working group), teaching contents and teaching methods (including the "teaching of ethics teaching"), research ethics vs. "ethics of research products", and ethical commitments to society vs. to clients. We also identified disciplines that are especially promising for our search, in particular medicine, bioethics, statistics, nuclear physics, psychology, journalism and legal sciences. Further relevant fields include data ethics, computer ethics, information ethics, robot ethics, and technology impact assessment. A number of stakeholders can be considered: digital right societies such as European Digital Rights (http://www.edri.org); civil rights societies, such as Transparency International (http://www.transparency.org); regulation authorities, such as the European Data Protection Supervisor (http://www.edps.eu); and professional associations, such as the Association of European Journalists (http://www.aej.org). There are also concepts such as dual use that have recurred throughout the Dagstuhl seminar and that have been studied in a large body of literature that we should become more familiar with. We discussed structural implementations of ethics, including ethics boards.

We share links and materials on a joint and open platform at https://etherpad.wikimedia.org/p/Dagstuhl_WG_Avoid_Reinventing.

10.3 Dynamics and Feedback in Discrimination Processes

Krishna P. Gummadi (MPI-SWS – Saarbrücken, DE), Chaitanya Baru (NSF – Arlington, US), Marina Drosou (Hellenic Police – Athens, GR), Salvatore Ruggieri (University of Pisa, IT), Rishiraj Saha Roy (MPI für Informatik – Saarbrücken, DE), Jannik Strötgen (MPI für Informatik – Saarbrücken, DE), and Suresh Venkatasubramanian (University of Utah – Salt Lake City, US)

In this breakout session, we identified different types of temporal aspects that are crucial in discrimination processes: (i) stepwise fairness, (ii) changing targets, and (iii) updated data. Thus, we propose a stepwise fairness model, where there is a final fairness goal but it is infeasible to achieve it at once due to a disproportionate loss in perceived "utilty". We should also consider the case when the perfect fairness scenario is a moving target, and changes with time. Relevant modeling paradigms for incremental fairness are online learning (continuous re-learning with new decisions), reinforcement learning (updating the reward function towards an optimal policy) and cooperative game theory (agents and algorithms as players trying to achieve ideal fairness). The general idea requires that the decisioning algorithm should incorporate feedback, i.e., current decisions should influence future decisions, and that the fairness "score" should be a function of the current time. The segregation model as proposed by Baroni and Ruggieri (2015) can be a good starting point for the stepwise fairness model.

10.4 Principles for Accountable Algorithms

Nicholas Diakopoulos (University of Maryland – College Park, US) and Sorelle Friedler (Haverford College, US)

Joint work of Marcelo Arenas, Solon Barocas, Nicholas Diakopoulos, Sorelle Friedler, Michael Hay, Bill Howe, H. V. Jagadish, Kris Unsworth, Arnaud Sahuguet, Suresh Venkatasubramanian, Christo Wilson, Cong Yu, Bendert Zevenbergen

Automated decision making algorithms are now used throughout industry and government, underpinning many processes from dynamic pricing to employment practices to criminal sentencing. Given that such algorithmically informed decisions have the potential for significant societal impact, the goal of this document is to help developers and product managers design and implement algorithmic systems in publicly accountable ways. Accountability in this context includes an obligation to report, explain, or justify algorithmic decision-making as well as mitigate any negative social impacts or potential harms.

We begin by outlining seven equally important guiding principles that follow from this premise:

Algorithms and the data that drive them are designed and created by people – There is always a human ultimately responsible for decisions made or informed by an algorithm. "The algorithm did it" is not an acceptable excuse if algorithmic systems make mistakes or have undesired consequences, including from machine-learning processes.

1. Auditability: Enable interested third parties to probe, understand, and review the behavior of the algorithm through disclosure of information that enables monitoring, checking, or criticism, including through provision of detailed documentation, technically suitable APIs, and permissive terms of use.
2. Error and Uncertainty: Identify, log, and articulate sources of error and uncertainty throughout the algorithm and its data sources so that expected and worst case implications can be understood and inform mitigation procedures.
3. Explainability: Ensure that algorithmic decisions as well as any data driving those decisions can be explained to end-users and other stakeholders in non-technical terms.
4. Fairness: Ensure that algorithmic decisions do not create discriminatory or unjust impacts when comparing across different demographics (e.g., race, sex, etc).
5. Human Experimentation: Consider the ethics of human experimentation in advance, including potential harms to end-users or others impacted by the algorithm, mitigation strategies for undue risk, and disclosure protocols for potential harms.
6. Privacy: Protect users' privacy surrounding any decisions or data derived or inferred from information about them.

7. Responsibility: Make available externally visible avenues of redress for adverse individual or societal effects of an algorithmic decision system, and designate an internal role for the person who is responsible for the timely remedy of such issues.

10.5 Data, Responsibly: Business and Research Opportunities

Julia Stoyanovich (Drexel University – Philadelphia, US), Serge Abiteboul (ENS – Cachan, FR), Chaitanya Baru (NSF – Arlington, US), Sorelle Friedler (Haverford College, US), Krishna P. Gummadi (MPI-SWS – Saarbrücken, DE), Michael Hay (Colgate University – Hamilton, US), Bill Howe (University of Washington – Seattle, US), Benny Kimelfeld (Technion – Haifa, IL), Arnaud Sahuguet (Cornell Tech NYC, US), Eric Simon (SAP France, FR), Suresh Venkatasubramanian (University of Utah – Salt Lake City, US), and Gerhard Weikum (MPI für Informatik – Saarbrücken, DE)

During this session we tried to identify business and funding opportunities around Data, Responsibly. On the business side, we looked at the consumer (PIMs, risk advisor, education), the enterprise (RaaS for Responsibility as a Service, Affordable Machine Learning) and tools (for auditing and benchmarking, for metrics, leveraging secure hardware, leveraging new language-based approaches and formal methods. We also looked at the social good element and some out of the box ideas such as data poisoning (for the dark web) and identity swapping.

On the research side, we agreed that a common vocabulary is truly necessary for collaboration and advertising. We suggested to try a Grand Challenge-based approach. Some challenges we identified include: universal health vault, vision zero data management (systems that can resist system design flaws and user errors), inequity in education and also a chance to design the next generation Internet starting from a clean slate. We tried to formulate the overall mission as "the science of data, responsibly and its application".

10.6 Explaining Decisions

Julia Stoyanovich (Drexel University – Philadelphia, US), Chaitanya Baru (NSF – Arlington, US), Claudia Bauzer Medeiros (UNICAMP – Campinas, BR), Krishna P. Gummadi (MPI-SWS – Saarbrücken, DE), Bill Howe (University of Washington – Seattle, US), Arnaud Sahuguet (Cornell Tech NYC, US), Jan Van den Bussche (Hasselt University, BE), and Gerhard Weikum (MPI für Informatik – Saarbrücken, DE)

In this session, we examined what it means to have an algorithm explain its decisions. We first considered the possible audience for such explanations. Stakeholders we have identified include: researcher, developer, consumer, policy maker, politician, regulator, competitor, and auditor.

We then looked at the goal we want to achieve with the explanation, e.g., explaining how the algorithm works, explaining the decision itself, making the recipient not feel bad about the decision, and providing actionable information for the recipient to get a better chance at a positive outcome in the future. An explanation should have the following properties: (a) understandable, (b) manageable (by generator and recipient), (c) consistent, (d) sound, (e) complete, (f) minimal.

We identified two broad categories of decisions. For allocation decisions, a set of limited resources needs to be allocated among agents who have expressed preferences. Decisions for all agents are generated at the same time by the algorithm. For single decisions, the explanation will highly depend on the nature of the algorithm used.

We then considered the kinds of explanations that may be appropriate for different kinds of algorithmic processes. For cases where the decision is based on a set of rules, an explanation can be the set of rules that triggered it. For cases where the decision is based on a decision tree, an explanation can be the full tree or the path (from root to leaf) that lead to the decision. For other cases, an explanation should identify the most critical attributes that contributed to the decision. We also identified another form of explanation where a macro-view of the behavior of the algorithm is visualized for the user, e.g., shared rides company response times shown on a city map.

The group concluded that explaining algorithmic decisions is an important, interesting and a technically challenging area that warrants attention from the research community.

10.7 Data, Responsibly: Use Cases and Benchmarking

Suresh Venkatasubramanian (University of Utah – Salt Lake City, US), Claudia Bauzer Medeiros (UNICAMP – Campinas, BR), Gerald Friedland (ICSI – Berkeley, US), Irini Fundulaki (FORTH – Heraklion, GR), and Salvatore Ruggieri (University of Pisa, IT)

This working group started to compile a set of use cases related to Data, Responsibly. Having such a set of curated examples can be extremely helpful when writing papers, applying for funding or simply trying to convince people about the importance and timeliness of the topic. This is work in progress but we hope that this evolving dataset will consist of a list of documented examples, annotated with the various dimensions that describe Data, Responsibly.

Participants

- Serge Abiteboul
ENS – Cachan, FR
- Marcelo Arenas
Pontificia Universidad Catolica de Chile, CL
- Solon Barocas
Microsoft – New York, US
- Chaitanya Baru
NSF – Arlington, US
- Claudia Bauzer Medeiros
UNICAMP – Campinas, BR
- Bettina Berendt
KU Leuven, BE
- Claude Castelluccia
INRIA – Grenoble, FR
- Nicholas Diakopoulos
University of Maryland – College Park, US
- Marina Drosou
Hellenic Police – Athen, GR
- Gerald Friedland
ICSI – Berkeley, US
- Sorelle Friedler
Haverford College, US
- Irini Fundulaki
FORTH – Heraklion, GR
- Krishna P. Gummadi
MPI-SWS – Saarbrücken, DE
- Michael Hay
Colgate Univ.y – Hamilton, US
- Bill Howe
University of Washington – Seattle, US
- H. V. Jagadish
University of Michigan – Ann Arbor, US
- Benny Kimelfeld
Technion – Haifa, IL
- Amélie Marian
Rutgers Univ. – Piscataway, US
- Pauli Miettinen
MPI für Informatik – Saarbrücken, DE
- Gerome Miklau
University of Massachusetts – Amherst, US
- Wolfgang Nejdl
Leibniz Univ. Hannover, DE
- Benjamin Nguyen
INSA – Bourges, FR
- Evaggelia Pitoura
University of Ioannina, GR
- Salvatore Ruggieri
University of Pisa, IT
- Rishiraj Saha Roy
MPI für Informatik – Saarbrücken, DE
- Arnaud Sahuguet
Cornell Tech NYC, US
- Eric Simon
SAP France, FR
- Julia Stoyanovich
Drexel Univ. – Philadelphia, US
- Jannik Strötgen
MPI für Informatik – Saarbrücken, DE
- Fabian Suchanek
Télécom ParisTech, FR
- Kristene Unsworth
Drexel Univ. – Philadelphia, US
- Jan Van den Bussche
Hasselt University, BE
- Suresh Venkatasubramanian
University of Utah – Salt Lake City, US
- Agnès Voisard
FU Berlin, DE
- Nicholas Weaver
ICSI – Berkeley, US
- Gerhard Weikum
MPI für Informatik – Saarbrücken, DE
- Christo Wilson
Northeastern University – Boston, US
- Cong Yu
Google – New York, US
- Ben Zevenbergen
University of Oxford, GB