Report from Dagstuhl Seminar 17141

Probabilistic Methods in the Design and Analysis of Algorithms

Edited by

Bodo Manthey[1], Claire Mathieu[2], Heiko Röglin[3], and Eli Upfal[4]

1 **University of Twente, NL,** b.manthey@utwente.nl
2 **ENS – Paris, FR,** cmathieu@di.ens.fr
3 **Universität Bonn, DE,** roeglin@cs.uni-bonn.de
4 **Brown University – Providence, US,** eliezer_upfal@brown.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17141 "Probabilistic Methods in the Design and Analysis of Algorithms".

Probabilistic methods play a central role in theoretical computer science. They are a powerful and widely applied tool used, for example, for designing efficient randomized algorithms and for establishing various lower bounds in complexity theory. They also form the basis of frameworks like average-case and smoothed analysis, in which algorithms are analyzed beyond the classical worst-case perspective. The seminar was on probabilistic methods with a focus on the design and analysis of algorithms.

The seminar helped to consolidate the research and to foster collaborations among the researchers who use probabilistic methods in different areas of the design and analysis of algorithms.

Seminar April 2–7, 2017 – http://www.dagstuhl.de/17141
1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems, G.1.6 Optimization, G.2 Discrete Mathematics, G.3 Probability and Statistics
Keywords and phrases analysis of algorithms, average-case analysis, random graphs, randomized algorithms, smoothed analysis, sub-linear algorithms
Digital Object Identifier 10.4230/DagRep.7.4.1
Edited in cooperation with Stefan Klootwijk

1 Executive Summary

Bodo Manthey
Claire Mathieu
Heiko Röglin
Eli Upfal

Probabilistic methods play a central role in theoretical computer science. They are a powerful and widely applied tool used, for example, for designing efficient randomized algorithms and for establishing various lower bounds in complexity theory. They also form the basis of frameworks like average-case and smoothed analysis, in which algorithms are analyzed beyond the classical worst-case perspective. The seminar was on probabilistic methods with a focus on the design and analysis of algorithms.

Probabilistic methods are often used in algorithm analysis when worst-case analysis does not provide useful or empirically accurate results. For example, worst-case analysis suggests

that the simplex method is an exponential-time algorithm for linear programming, while in fact it runs in near-linear time on almost all inputs of interest. For the simplex method and many other algorithms worst-case inputs are often rather contrived and occur hardly ever in practical applications. The last decade has seen much interest in the development of a more realistic and robust algorithmic theory that is not entirely based on worst-case performance. One very successful line of research studies the performance of algorithms on inputs that are to some extent random. Besides average-case analysis, in which inputs are generated randomly according to some fixed distribution, also more sophisticated semi-random models have gained momentum.

Another area in which probabilistic methods play a central role is stochastic optimization. Here uncertainty in the data is modeled by probability distributions and the actual data is only revealed over time. For example, in a scheduling problem one might know the probability distribution of a job's length but one learns its actual length only by executing it.

Probabilistic methods are also central in algorithm design. For many optimization problems, the most efficient known algorithms rely essentially on randomization. In other areas, like sublinear algorithms and hashing, one can even prove that randomization is necessary to obtain good algorithms.

The seminar covered recent progress in the context of probabilistic methods.

2 Table of Contents

3 Introduction

Average-case Analysis. When worst-case analysis fails to predict the performance of an algorithm, average-case analysis is often the first alternative. Average-case analysis has been the subject of a huge amount of research in the last decades, building on insights from, for instance, random graph theory. However, often very simple probability distributions, such as Erdős-Rényi random graphs, have been used in order to keep the average-case analysis tractable. Only recently, average-case analysis has adopted more complex input models, inspired by the study of complex networks. Examples are the usage of first-passage percolation to generate random metric spaces (Bringmann et al., *Algorithmica*, 2015), incorporating gender in social network models to model the glass ceiling effect (Avin et al., *ITCS* 2015), and the preferential attachment process to generate random graphs whose node degrees follow a power-law distribution.

Smoothed Analysis. The drawback of average-case analysis is that purely random instances often have very special properties that distinguish them from typical instances. In smoothed analysis one considers adversarial inputs that are slightly perturbed at random. Spielman and Teng, who invented this notion in 2001, considered adversarial linear programs in which every coefficient is perturbed by Gaussian noise with standard deviation σ. They proved that the expected running time of a certain variant of the simplex algorithm on such linear programs is polynomially bounded in the input size and $1/\sigma$. Hence, even if the amount of randomness is small, the expected running time of the simplex method is polynomially bounded.

Smoothed analysis can be viewed as a less pessimistic worst-case analysis, in which the randomness rules out pathological worst-case instances that are rarely observed in practice but dominate the worst-case analysis. After its invention it has attracted a great deal of attention and it has been applied in a variety of different contexts, e.g., in multi-objective optimization, local search, clustering, and online algorithms.

Semi-random Graph Models. Semi-random input models can be considered as analogues of smoothed analysis for graph problems and they even predate smoothed analysis by a couple of years. There is a variety of semi-random graph models that go beyond the classical Erdős-Rényi random graphs. In most of these models graphs are generated by a noisy adversary – an adversary whose decisions (whether or not to insert a particular edge) have some small probability of being reversed. Another well-studied class of semi-random models are *planted models*, in which a solution (e.g., an independent set or a partitioning of the vertices in color classes) is chosen and then edges are added randomly or by an adversary of limited power in such a way that the given solution stays a valid solution for the given problem.

Similar to smoothed analysis, semi-random models have been invented in order to better understand the complexity of NP-hard graph problems because Erdős-Rényi random graphs often do not reflect the instances one encounters in practice – many graph problems are efficiently solvable on such random graphs.

Online Problems with Random Order. The competitive ratio approach to evaluating online algorithms can be separated into two parts; the sequence of items given as input to an online problem can form a worst-case set of items; and the order in which they are revealed can be a worst-case order. An intermediate measure, between worst-case competitive ratio and average-case competitive ratio when each item in the sequence is drawn independently from some distribution, occurs when the set of input items is worst-case but the arrival order is

random, as might happen if the requests come from sources that do not coordinate their timing. Many problems can be solved more efficiently when the arrival order is random.

Algorithm Design based on Probabilistic Analysis. It is an interesting question if insights from probabilistic analyses of algorithms can give rise to improved algorithms. Recently, some steps in this direction have been taken. The smoothed analysis of the shadow vertex simplex method has, for example, led to efficient algorithms to find short paths on certain polyhedra and improved diameter bounds (Dadush, Hähnle, *SoCG*, 2015).

Stochastic Optimization. In stochastic optimization one has to cope with uncertainty in the inputs, which is usually modeled by probability distributions. One illustrative example are scheduling problems in which the processing times of the tasks follow known probability distributions and the actual processing time of a job is not known until it is completed. In recent years, there has been significant progress in stochastic optimization, in particular, in the area of stochastic scheduling.

Sample Complexity and Sublinear Algorithms. Sampling is a powerful technique, which is at the core of many randomization techniques. Using a finite, often small, set of observations, we attempt to estimate properties of an entire sample space. How good are estimates obtained from a sample? Any rigorous application of sampling requires an understanding of the *sample complexity* of the problem – the minimum size sample needed to obtain the required results. The problem becomes particularly difficult when we try to use one sample set (one set of observations) to detect, or estimate the probabilities of a large (possibly infinite) set of events. Sample complexity has been extensively studied in the context of statistical learning theory, however recent applications have shown the importance of these issues in related fields such as data mining and big data analysis. In particular, there has been significant interest in applying Rademacher averages to improve progressive sampling, and VC-dimension to the study of social network properties.

Related to sample complexity is the area of sublinear algorithms. For huge data sets, even linear-time algorithms are impractical. Hence, the area of sublinear algorithms has developed. Sublinear algorithms inherently need randomness to sample from the input.

Program of the Seminar

The program of the seminar consisted of 29 talks, including the following survey talks:

- Secretary problems and prophet inequalities, Thomas Kesselheim;
- New (and old) results about the Moran Process, Leslie Ann Goldberg;
- Randomized shared memory algorithms, Philipp Wölfel;
- Rademacher averages, Matteo Riondata.

The rest of the talks were 25-minute presentations on recent research of the participants. The time between lunch and the afternoon coffee break was mostly left open for individual discussions and collaborations in small groups. One of the main goals of the seminar was to foster collaborations among the researchers using probabilistic methods in different ways in the area of algorithms research. The feedback provided by the participants shows that the goals of the seminar, namely to circulate new ideas and create new collaborations, were met to a large extent. The organizers and participants wish to thank the staff and the management of Schloss Dagstuhl for their assistance and support in the arrangement of a very successful and productive event.

4 Overview of Talks

4.1 Stochastic Control via Entropy Compression

Dimitris Achlioptas (University of California – Santa Cruz, US)

Joint work of Dimitris Achlioptas, Fotis Iliopoulos, Nikos Vlassis
Main reference D. Achlioptas, F. Iliopoulos, N. Vlassis, "Stochastic Control via Entropy Compression", in Proc. of the 44th International Colloquium on Automata, Languages, and Programming (ICALP 2017), LIPIcs, Vol 80, pp. 83:1–83:13, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL http://dx.doi.org/10.4230/LIPIcs.ICALP.2017.83

We consider an agent trying to bring a system to an acceptable state by repeated probabilistic action (stochastic control). Specifically, in each step the agent observes the flaws present in the current state, selects one of them, and addresses it by probabilistically moving to a new state, one where the addressed flaw is most likely absent, but where one or more new flaws may be present. Several recent works on algorithmizations of the Lovász Local Lemma (LLL) have established sufficient conditions for such an agent to succeed. Motivated by the paradigm of Partially Observable Markov Decision Processes (POMDPs) we study whether such stochastic control is also possible in a noisy environment, where both the process of state-observation and the process of state-evolution are subject to adversarial perturbation (noise). The introduction of noise causes the tools developed for LLL algorithmization to break down since the key LLL ingredient, the sparsity of the causality (dependence) relationship, no longer holds. To overcome this challenge we develop a new analysis where entropy plays a central role, both to measure the rate at which progress towards an acceptable state is made and the rate at which the noise undoes this progress. The end result is a sufficient condition that allows a smooth trade-off between the intensity of the noise and the amenability of the system, recovering an asymmetric LLL condition in the noiseless case. To our knowledge, this is the first tractability result for a nontrivial class of POMDPs under stochastic memoryless control.

4.2 Friend or Foe? Population Protocols can perform Community Detection

Luca Becchetti (Sapienza University of Rome, IT)

Joint work of Luca Becchetti, Andrea Clementi, Emanuele Natale, Francesco Pasquale, Prasad Raghavendra, Luca Trevisan
Main reference L. Becchetti, A. E. F. Clementi, E. Natale, F. Pasquale, P. Raghavendra, L. Trevisan, "Friend or Foe? Population Protocols can perform Community Detection", arXiv:1703.05045 [cs.DM], 2017.
URL https://arxiv.org/abs/1703.05045

We present a simple distributed algorithm that, given a regular graph consisting of two communities (or clusters), each inducing a good expander and such that the cut between them has sparsity polylogarithmic recovers the two communities.

More precisely, upon running the protocol, every node assigns itself a binary label of logarithmic length m, so that with high probability, for all but a small number of outliers, nodes within the same community are assigned labels with Hamming distance $o(m)$, while nodes belonging to different communities receive labels with Hamming distance at least $m/2 - o(m)$. We refer to such an outcome as a *community sensitive labeling* of the graph.

Our algorithm uses polylogarithmic local memory and computes the community sensitive labeling after each node performs polylogarithmic steps of local work.

Our algorithm and its analysis work in the *(random) population protocol* model, in which anonymous nodes do not share any global clock (the model is asynchronous) and communication occurs over one single (random) edge per round. We believe, this is the first provably-effective protocol for community detection that works in this model.

4.3 Ignore or Comply? On Breaking Symmetry in Consensus

Petra Berenbrink (Universität Hamburg, DE)

Joint work of Petra Berenbrink, Andrea Clementi, Robert Elsässer, Peter Kling, Frederik Mallmann-Trenn, Emanuele Natale
Main reference P. Berenbrink, A. E. F. Clementi, R. Elsässer, P. Kling, F. Mallmann-Trenn, E. Natale, "Ignore or Comply? On Breaking Symmetry in Consensus", to appear in PODC 2017.
URL https://arxiv.org/abs/1702.04921

We study consensus processes on the complete graph of n nodes. Initially, each node supports one from up to n opinions. Nodes randomly and in parallel sample the opinions of constant many nodes. Based on these samples, they use an update rule to change their own opinion. The goal is to reach consensus, a configuration where all nodes support the same opinion. We compare two well-known update rules: 2-Choices and 3-Majority. In the former, each node samples two nodes and adopts their opinion if they agree. In the latter, each node samples three nodes: If an opinion is supported by at least two samples the node adopts it, otherwise it randomly adopts one of the sampled opinions. Known results for these update rules focus on initial configurations with a limited number of colors, or typically assume a bias, where one opinion has a much larger support than any other. For such biased configurations, the time to reach consensus is roughly the same for 2-Choices and 3-Majority. Interestingly, we prove that this is no longer true for configurations with a large number of initial colors. In particular, we show that 3-Majority reaches consensus with high probability in $O(n^{3/4} \log^{7/8} n)$ rounds, while 2-Choices can need $\Omega(n/\log n)$ rounds. We thus get the first unconditional sublinear bound for 3-Majority and the first result separating the consensus time of these processes. Along the way, we develop a framework that allows a fine-grained comparison between consensus processes from a specific class. We believe that this framework might help to classify the performance of more consensus processes.

4.4 Information-theoretic thresholds

Amin Coja-Oghlan (Goethe-Universität – Frankfurt am Main, DE)

Joint work of Amin Coja-Oghlan, Florent Krzakala, Will Perkins, Lenka Zdeborová
Main reference A. Coja-Oghlan, F. Krzakala, W. Perkins, L. Zdeborová, "Information-theoretic thresholds from the cavity method", arXiv:1611.00814 [cs.DM], 2016.
URL https://arxiv.org/abs/1611.00814

Vindicating a sophisticated but non-rigorous physics approach called the cavity method, we establish a formula for the mutual information in statistical inference problems induced by random graphs and we show that the mutual information holds the key to understanding

certain important phase transitions in random graph models. We work out several concrete applications of these general results. For instance, we pinpoint the exact condensation phase transition in the Potts antiferromagnet on the random graph, thereby improving prior approximate results [1]. Further, we prove the conjecture from [3] about the condensation phase transition in the random graph coloring problem for any number $q \geq 3$ of colors. Moreover, we prove the conjecture on the information-theoretic threshold in the disassortative stochastic block model [2]. Additionally, our general result implies the conjectured formula for the mutual information in Low-Density Generator Matrix codes [4].

References

1 P. Contucci, S. Dommers, C. Giardinà, S. Starr. *Antiferromagnetic Potts model on the Erdős-Rényi random graph.* Communications in Mathematical Physics 323 (2013) 517–554.
2 A. Decelle, F. Krzakala, C. Moore, L. Zdeborová. *Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications.* Phys. Rev. E 84 (2011) 066106.
3 F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová. *Gibbs states and the set of solutions of random constraint satisfaction problems.* Proc. National Academy of Sciences 104 (2007) 10318–10323.
4 A. Montanari. *Tight bounds for LDPC and LDGM codes under MAP decoding.* IEEE Transactions on Information Theory 51 (2005) 3221-3246.

4.5 Making Banaszczyk's Bound Constructive for the Komlos Problem

Daniel Dadush (CWI – Amsterdam, NL)

Joint work of Nikhil Bansal, Daniel Dadush, Shashwat Garg

We consider the problem of finding a low discrepancy coloring for sparse set systems where each element lies in at most t sets. We give an efficient algorithm that finds a coloring with discrepancy $O((t \log n)^{1/2})$, matching the best known non-constructive bound for the problem due to Banaszczyk. The previous algorithms only achieved an $O(t^{1/2} \log n)$ bound. The result also extends to the more general Komlos setting, where each vector has norm at most 1, and gives an algorithmic $O(\log^{1/2} n)$ bound.

Joint work with Nikhil Bansal and Shashwat Garg.

4.6 Dual-Pivot Quicksort: Optimality, Analysis and Zeros of Lattice Paths

Martin Dietzfelbinger (TU Ilmenau, DE)

Joint work of Martin Aumüller, Martin Dietzfelbinger, Daniel Krenn, Clemens Heuberger, Helmut Prodinger
Main reference M. Aumüller, M. Dietzfelbinger, C. Heuberger, D. Krenn, H. Prodinger, "Dual-Pivot Quicksort: Optimality, Analysis and Zeros of Associated Lattice Paths", arXiv:1611.00258 [math.CO], 2016.
URL https://arxiv.org/abs/1611.00258

Quicksort is a venerable sorting algorithm, it is taught in basic algorithms classes, and it is routinely used in practice. Can there be anything new about Quicksort today? *Dual-pivot*

quicksort refers to variants of classical quicksort where in the partitioning step two pivots are used to split the input into three segments. Algorithms of this type had been studied by Sedgewick (1975) and by Hennequin, but they received new attention starting from 2009, when a dual-pivot algorithm due to Yaroslavskiy, Bentley, and Bloch replaced the well-engineered quicksort algorithm in Oracle's Java 7 runtime library. An analysis of a simplified variant of this algorithm by Nebel and Wild from 2012 gave $1.9 \ln n$ comparisons on average for n input numbers. (Other works ensued. Standard quicksort has $2n \ln n$ expected comparisons.) In the center of the analysis is the partitioning procedure, which splits the input keys in "small" (smaller than small pivot), "medium" (between the two pivots), "large" (larger than large pivot). We identify a partitioning strategy with the minimum average number of key comparisons. It keeps count of how many large and small elements were seen before and prefers the corresponding pivot. The comparison count is closely related to a "random walk" on the integers which keeps track of the difference of large and small elements seen so far. (An alternative way of understanding what is going on is a Pólya urn with three colors.) For the fine analysis it is essential to understand the expected number of times this random walk hits zero. The expected number of comparisons can be determined exactly and as a formula up to lower terms. It is $1.8n \ln n - 2.38...n + 1.675 \ln n + O(1)$.

4.7 Randomized Rumor Spreading Revisited

Benjamin Doerr (Ecole Polytechnique – Palaiseau, FR)

We develop a simple and generic method to analyze randomized rumor spreading processes in fully connected networks. In contrast to all previous works, which heavily exploit the precise definition of the process under investigation, we only need to understand the probability and the covariance of the events that uninformed nodes become informed. This universality allows us to easily analyze the classic push, pull, and push-pull protocols both in their pure version and in several variations such as messages failing with constant probability or nodes calling a random number of others each round. Some dynamic models can be analyzed as well, e.g., when the network is a $G(n, p)$ random graph sampled independently each round [Clementi et al. (ESA 2013)].

Despite this generality, our method determines the expected rumor spreading time precisely apart from additive constants, which is more precise than almost all previous works. We also prove tail bounds showing that a deviation from the expectation by more than an additive number of r rounds occurs with probability at most $\exp(-\Omega(r))$.

We further use our method to discuss the common assumption that nodes can answer any number of incoming calls. We observe that the restriction that only one call can be answered leads to a significant increase of the runtime of the push-pull protocol. In particular, the double logarithmic end phase of the process now takes logarithmic time. This also increases the message complexity from the asymptotically optimal $\Theta(n \log \log n)$ [Karp, Shenker, Schindelhauer, Vöcking (FOCS 2000)] to $\Theta(n \log n)$. We propose a simple variation of the push-pull protocol that reverts back to the double logarithmic end phase and thus to the $\Theta(n \log \log n)$ message complexity.

4.8 Thompson Sampling For Stochastic Bandits with Graph Feedback

Devdatt Dubhashi (Chalmers UT – Göteborg, SE)

Joint work of Aristide Tossou, Christos Dimitrakakis, Devdatt Dubhashi
Main reference A. C. Y. Tossou, C. Dimitrakakis, D. P. Dubhashi, "Thompson Sampling for Stochastic Bandits with Graph Feedback", in Proc. of the 31st AAAI Conf. on Artificial Intelligence (AAAI 2017), pp. 2660-2666, 2017.
URL https://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/14768

We present a novel extension of Thompson Sampling for stochastic sequential decision problems with graph feedback, even when the graph structure itself is unknown and/or changing. We provide theoretical guarantees on the Bayesian regret of the algorithm, linking its performance to the underlying properties of the graph. Thompson Sampling has the advantage of being applicable without the need to construct complicated upper confidence bounds for different problems. We illustrate its performance through extensive experimental results on real and simulated networks with graph feedback. More specifically, we tested our algorithms on power law, planted partitions and Erdős-Rényi graphs, as well as on graphs derived from Facebook and Flixster data. These all show that our algorithms clearly outperform related methods that employ upper confidence bounds, even if the latter use more information about the graph.

4.9 New (and old) results about the Moran Process

Leslie Ann Goldberg (University of Oxford, GB)

Joint work of Andreas Galanis, Andreas Goebel, Leslie Ann Goldberg, John Lapinskas, Johannes Lengler, Florian Meier, Konstantinos Panagiotou, Pascal Pfister, David Richerby
Main reference L. A. Goldberg, J. Lapinskas, J. Lengler, F. Meier, K. Panagiotou, P. Pfister, "Asymptotically Optimal Amplifiers for the Moran Process", arXiv:1611.04209 [math.PR], 2016.
URL https://arxiv.org/abs/1611.04209
Main reference A. Galanis, A. Goebel, L. A. Goldberg, J. Lapinskas, D. Richerby, "Amplifiers for the Moran Process", Journal of the ACM (JACM) 64(1):5, 2017.
URL http://dx.doi.org/10.1145/3019609

The Moran process, as adapted by Lieberman, Hauert and Nowak, is a model of a population on a graph, evolving in discrete time. Individuals in the population are associated with the vertices of the graph. Certain individuals, called "mutants" have fitness r and other individuals, called "non-mutants" have fitness 1. The state of being a mutant or not can be spread from vertices to neighbours. We focus on the situation where the mutation is advantageous, in the sense that $r > 1$. If the graph is strongly connected then, with probability 1, the process will either reach the state where there are only mutants (known as fixation) or it will reach the state where there are only non-mutants (known as extinction). A set of (directed or undirected) graphs is said to be strongly amplifying if the extinction probability tends to 0 when the Moran process is run on graphs in this set, starting from the state with a single mutation, at a uniformly-chosen vertex. It turns out that strong amplifiers exist, even in the undirected case. This talk will tell you what is known about them, including joint work with Galanis, Goebel, Lapinskas and Richerby and also joint work with Lapinskas, Lengler, Meier, Panagiotou and Pfister.

4.10 Resilience of the chromatic number

Ross Kang (Radboud University Nijmegen, NL)

Joint work of Ross Kang, Colin McDiarmid

For p bounded away from 1 and $np \to \infty$, the asymptotic value of the chromatic number of the binomial random graph $G(n,p)$ is resilient to the removal of a subgraph H with $o(n\log(np))$ edges. This is best possible. Our work is counterpoint to work of Sudakov and Vu (2008) and Alon and Sudakov (2010), which was on chromatic resilience for $G(n,p)$ subject to the addition, rather than removal, of edges. Joint work with Colin McDiarmid.

4.11 Matching Prophets and Secretaries

Thomas Kesselheim (TU Dortmund, DE)

Main reference T. Kesselheim, K. Radke, A. Tönnis, B. Vöcking, "An Optimal Online Algorithm for Weighted Bipartite Matching and Extensions to Combinatorial Auctions", Algorithms – ESA 2013, LNCS, Vol. 8125, pp. 589–600, Springer, 2013.
URL http://dx.doi.org/10.1007/978-3-642-40450-4_50
Main reference M. Feldman, N. Gravin, B. Lucier, "Combinatorial Auctions via Posted Prices", in Proc. of the 26th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA 2015), pp. 123–135, 2015.
URL http://dx.doi.org/10.1137/1.9781611973730.10

We discuss two different approaches for probabilistic input models in online settings. In the first model, the optimization instance is determined by an adversary but it is revealed in random order. The secretary problem is a common representative of this class. In the second model, the optimization instance is drawn from known probability distributions and presented in adversarial order, like for the prophet inequality. For both settings, we survey different online problems with a focus on edge-weighted bipartite matching, for which optimal guarantees are known.

4.12 Probabilistic Analysis of Facility Location on Random Shortest Path Metrics

Stefan Klootwijk (University of Twente, NL)

Joint work of Stefan Klootwijk, Bodo Manthey
Main reference S. Klootwijk, "Probabilistic Analysis of Facility Location on Random Shortest Path Metrics", Master Thesis, University of Twente, 2016.
URL http://purl.utwente.nl/essays/71002

We investigate the performance of the facility location problem using the model of random shortest path metrics. We analyze some probabilistic properties for a simple heuristic which provides a solution to the facility location problem: opening a certain number of arbitrary facilities (with that certain number only depending on the facility opening cost). We show that, for almost any facility opening cost, this heuristic yields a $1 + o(1)$ approximation in expectation. In the remaining few cases we show that this heuristic yields an $O(1)$ approximation in expectation.

4.13 Towards $(1+\varepsilon)$-Approximate Flow Sparsifiers

Robert Krauthgamer (Weizmann Institute – Rehovot, IL)

Joint work of Alexandr Andoni, Anupam Gupta, Robert Krauthgamer
Main reference A. Andoni, A. Gupta, R. Krauthgamer, "Towards $(1+\varepsilon)$-Approximate Flow Sparsifiers", in Proc. of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2014), pp. 279–293, 2014.
URL http://dx.doi.org/10.1137/1.9781611973402.20

A key challenge in designing graph algorithms is to "compress" a graph G so as to preserve some of its basic properties, such as distances and cuts. I will discuss vertex-sparsification, where given a graph G with k terminal vertices, the goal is to construct a small graph G' containing the same terminals, such that all cuts/flows between the terminals in G' approximate those in G within factor $1+\varepsilon$.

I will present a randomized construction for bipartite graphs, and discuss some open questions.

4.14 Algorithms and mechanisms for two-sided bayesian auctions.

Stefano Leonardi (Sapienza University of Rome, IT)

Joint work of Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Stefano Leonardi, Tim Roughgarden, Stefano Turchetta
Main reference R. Colini-Baldeschi, B. de Keijzer, S. Leonardi, S. Turchetta, "Approximately Efficient Double Auctions with Strong Budget Balance", in Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), pp. 1424–1443, 2016.
URL https://doi.org/10.1137/1.9781611974331.ch98
Main reference R. Colini-Baldeschi, P. W. Goldberg, B. de Keijzer, S. Leonardi, T. Roughgarden, S. Turchetta, "Approximately Efficient Two-Sided Combinatorial Auctions", in Proc. of the 2017 ACM Conf. on Economics and Computation (EC'17), pp. 591–608, ACM, 2017.
URL https://doi.org/10.1145/3033274.3085128

Mechanism design for one-sided markets is an area of extensive research in economics and, since more than a decade, in computer science as well. Two-sided markets, on the other hand, have not received the same attention despite the numerous applications to Internet advertisement. We present algorithms and mechanisms for two-sided markets in which both buyers and sellers act strategically.

An ideal goal in two-sided markets is to maximize the social welfare of buyers and sellers with individually rational (IR), incentive compatible (IC) and budget-balanced mechanisms (BB), which requires that the mechanism does not subsidize the market or make an arbitrary profit from the exchange. Unfortunately, Myerson and Satterthwaite proved in 1983 that this goal cannot be achieved even in the bayesian setting and for the simple case of only one buyer and one seller. In this talk I'll discuss meaningful trade-offs and algorithmic approximations of the above requirements for the bayesian setting.

I'll present simple two-sided sequential posted price mechanisms which provides an $O(1)$ approximation to the optimal social welfare while obeying the IR, IC and BB constraints. The mechanisms work for any number of buyers and sellers with arbitrary, independent distributions and matroid constraints for the case of unit supply buyers and sellers. These mechanisms can be extended to two-sided combinatorial auctions with additive and XOS valuations on subsets of items.

4.15 Distance in the Forest Fire Model – How far are you from Eve?

Frederik Mallmann-Trenn (ENS – Paris, FR)

Joint work of Claire Mathieu, Frederik Mallmann-Trenn, Varun Kanade, Reut Levi, Zvi Lotker
Main reference V. Kanade, R. Levi, Z. Lotker, F. Mallmann-Trenn, C. Mathieu, "Distance in the Forest Fire Model How far are you from Eve?", in Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), pp. 1602–1620, 2016.
URL http://dx.doi.org/10.1137/1.9781611974331.ch109

Leskovec, Kleinberg and Faloutsos (2005) observed that many social networks exhibit properties such as shrinking (i.e. bounded) diameter, densification, and (power-law) heavy tail degree distributions. To explain these phenomena, they introduced a generative model, called the Forest Fire model, and using simulations showed that this model indeed exhibited these properties; however, proving this rigorously was left as an open problem.

In this paper, we analyse one of these properties, shrinking diameter. We define a restricted version of their model that incorporates the main features that seem to contribute towards this property, and prove that the graphs generated by this model exhibit shrinking distance to the seed graph. We prove that an even simpler model, the random walk model, already exhibits this phenomenon.

4.16 The noisy skyline problem

Claire Mathieu (ENS – Paris, FR)

Joint work of Frederik Mallmann-Trenn, Claire Mathieu

Given a set of points in a d-dimensional space, we seek to compute the skyline, i.e., those points which are not strictly dominated by any other point, using few comparisons between elements. We study the setting where comparisons can fail with constant probability. This version of the problem was introduced and studied by Groz & Milo. We provide two algorithms computing the skyline with query complexity $O(nd\log(dk))$ and $O(ndk\log(k))$, where k is the size of the skyline. These results improve on the bounds of Groz and Milo by a factor of k. Furthermore, we prove that the obtained bounds are (up to constants) tight when d is not too large.

4.17 Advances in Streaming Matching

Morteza Monemizadeh (Charles University – Prague, CZ)

Joint work of Rajesh Chitnis, Graham Cormode, Hossein Esfandiari, Elena Grigorescu, Mohammad Taghi Hajiaghayi, Hossein Jowhari, Vahid Liaghat, Andrew McGregor, Morteza Monemizadeh, S. Muthukrishnan, Krzysztof Onak, Sofya Vorotnikova, Samson Zhou

Main reference G. Cormode, H. Jowhari, M. Monemizadeh, S. Muthukrishnan, "The Sparse Awakens: Streaming Algorithms for Matching Size Estimation in Sparse Graphs", in Proc. of the 25th Annual European Symposium on Algorithms (ESA 2017), LIPIcs, Vol. 87, pp. 29:1–29:15, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.

URL https://arxiv.org/abs/1608.03118

Main reference R.H. Chitnis, G. Cormode, H. Esfandiari, M. T. Hajiaghayi, A. McGregor, M. Monemizadeh, S. Vorotnikova, "Kernelization via Sampling with Applications to Finding Matchings and Related Problems in Dynamic Graph Streams", in Proc. of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2016), pp. 1326–1344, 2016.

URL http://dx.doi.org/10.1137/1.9781611974331.ch92

Main reference R.H. Chitnis, G. Cormode, M.T. Hajiaghayi, M. Monemizadeh, "Parameterized Streaming: Maximal Matching and Vertex Cover", in Proc. of the 26th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA 2015), pp. 1234–1251, 2015.

URL http://dx.doi.org/10.1137/1.9781611973730.82

Main reference H. Esfandiari, M.T. Hajiaghayi, V. Liaghat, M. Monemizadeh, K. Onak, "Streaming Algorithms for Estimating the Matching Size in Planar Graphs and Beyond", in Proc. of the 26th Annual ACM-SIAM Symp. on Discrete Algorithms (SODA 2015), pp. 1217–1233, 2015.

URL http://dx.doi.org/10.1137/1.9781611973730.81

As noted by Lovász and Plummer in their classic book, "Matching Theory is a central part of graph theory, not only because of its applications, but also because it is the source of important ideas developed during the rapid growth of combinatorics during the last several decades."

We consider variants of matching in data streams when the insertion and the deletion of edges are revealed in a streaming fashion. In particular,

1. When the input graph is planar, we present a simple and elegant streaming algorithm that with high probability estimates the size of a maximum matching within a constant factor using $O(n^{2/3})$ space, where n is the number of vertices. The approach generalizes to the family of graphs that have bounded arboricity, which include graphs with an excluded constant-size minor.
2. We further reduce the required memory size to $O(\sqrt{n})$ for three restricted settings: (i) when the input graph is a forest; (ii) when we have 2-passes and the input graph has bounded arboricity; and (iii) when the edges arrive in random order and the input graph has bounded arboricity.
3. We present a simple but powerful subgraph sampling primitive that is applicable in a variety of computational models including dynamic graph streams (where the input graph is defined by a sequence of edge/hyperedge insertions and deletions) and distributed systems such as MapReduce. In the case of dynamic graph streams, we use this primitive to prove that there exists an $O(k^2)$ space algorithm that returns the edges of a maximum matching on the assumption the cardinality is at most k. The best previous algorithm used $O(kn)$ space. We prove our result is optimal up to logarithmic factors. Our algorithm has $O(1)$ update time.
4. We also show that there exists an $O(n^2/\alpha^3)$ space algorithm that returns an α-approximation for matchings of arbitrary size. In independent work, Assadi et al. (SODA 2016) proved this approximation algorithm is optimal and provided an alternative algorithm. We generalize our exact and approximate algorithms to weighted matching. For graphs with low arboricity such as planar graphs, the space required for constant approximation can be further reduced. While there has been a substantial amount of work on approximate matching in insert-only graph streams, these are the first non-trivial results in the dynamic setting.

4.18 Multi-label classification with pairwise relations

Seffi Naor (Technion – Haifa, IL)

Joint work of Shahar Chen, Dotan Di Castro, Zohar Karnin, Liane Lewin-Eytan, Seffi Naor, Roy Schwartz

Motivated by various applications in machine learning, e.g., classification of web pages, semantic tagging of images, and functional genomics, we introduce the metric multi-labeling problem. The input is a set of objects with pairwise relations, given as a weighted graph over the objects, together with a label set and assignment costs of labels to objects. Departing from classical work on the metric labeling problem, assignment costs may be either positive or negative, reflecting a recommendation given by a local learning process which infers label preferences of objects. However, the learning process ignores pairwise relations of objects. Clearly, the labeling cost for completely agreeing with this recommendation is the minimum possible, and this is our benchmark labeling. The objective is to find an assignment of (multiple) labels to objects, at least one for each object, minimizing the sum of the following two terms: (1) the deviation cost from the benchmark labeling, and (2) the separation cost between pairs of similar objects receiving different label sets. Thus, the goal is to globally optimize the assignment of label sets to objects, while taking into account local preferences of objects. We obtain a (tight) 2-approximation for this problem.

Additionally, in many natural applications a bound is given on the number of labels that can be assigned to an object. This leads to a very interesting problem of simultaneous rounding of many (fractional) points contained in a uniform matroid polytope, while bounding pairwise distances. We achieve a true $O(\log k)$-approximation for this problem (k is the number of labels).

4.19 The distortion of Locality Sensitive Hashing

Alessandro Panconesi (Sapienza University of Rome, IT)

Joint work of Flavio Chierichetti, Ravi Kumar, Alessandro Panconesi, Erisa Terolli
Main reference F. Chierichetti, R. Kumar, A. Panconesi, E. Terolli, "The Distortion of Locality Sensitive Hashing", in Proc. of the 8th Innovations in Theoretical Computer Science Conf. (ITCS 2017), LIPIcs, Vol. 67, pp. 54:1–54:18, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
URL http://dx.doi.org/10.4230/LIPIcs.ITCS.2017.54

Given a pairwise similarity notion between objects, locality sensitive hashing (LSH) aims to construct a hash function family over the universe of objects such that the probability two objects hash to the same value is their similarity. LSH is a powerful algorithmic tool for large-scale applications and much work has been done to understand LSHable similarities, i.e., similarities that admit an LSH.

In this paper we focus on similarities that are provably non-LSHable and propose a notion of distortion to capture the approximation of such a similarity by a similarity that is LSHable. We consider several well-known non-LSHable similarities and show tight upper and lower bounds on their distortion.

4.20 MapReduce and Streaming Algorithms for Diversity Maximization in Metric Spaces of Bounded Doubling Dimension

Andrea Pietracaprina (University of Padova, IT)

Joint work of Matteo Ceccarello, Andrea Pietracaprina, Geppino Pucci, Eli Upfal
Main reference M. Ceccarello, A. Pietracaprina, G. Pucci, E. Upfal, "MapReduce and Streaming Algorithms for Diversity Maximization in Metric Spaces of Bounded Doubling Dimension", PVLDB 10(5):469–480, 2017.
URL http://doi.org/10.14778/3055540.3055541

Given a dataset of points in a metric space and an integer k, a diversity maximization problem requires determining a subset of k points maximizing some diversity objective measure (e.g., the minimum or the average distance between a pair of points in the subset). Diversity maximization problems are computationally hard, hence only approximate solutions can be hoped for. Although its applications are mostly in massive data analysis, most of the past research on diversity maximization has concentrated on the standard sequential setting. In this talk, we present novel space and pass/round-efficient approximation algorithms for diversity maximization in the Streaming and MapReduce models, which can handle very large inputs. We show that for point sets belonging to metric spaces of bounded doubling dimension, our algorithms attain an $(\alpha + \varepsilon)$-approximation ratio, for any constant $\varepsilon > 0$, where α is the best approximation ratio achieved by a polynomial-time, linear-space sequential algorithm. We also provide extensive experimental evidence of the effectiveness of our algorithms on both real world and synthetic datasets, scaling up to over a billion points. While our algorithms are deterministic, we show how randomization can help attain lower space requirements, and believe this to be a promising direction for future improvements. Also, we will briefly discuss recent results on the generalization of our approach to diversity maximization problems under matroid constraints.

4.21 Efficient Diameter Approximation for Large Graphs in MapReduce

Geppino Pucci (University of Padova, IT)

Joint work of Matteo Ceccarello, Andrea Pietracaprina, Geppino Pucci, Eli Upfal
Main reference M. Ceccarello, A. Pietracaprina, G. Pucci, E. Upfal, "Space and Time Efficient Parallel Graph Decomposition, Clustering, and Diameter Approximation", ACM SPAA 2015: 182–191.
Main reference M. Ceccarello, A. Pietracaprina, G. Pucci, E. Upfal, "A Practical Parallel Algorithm for Diameter Approximation of Massive Weighted Graphs", IEEE IPDPS 2016: 12–21.
URL http://dx.doi.org/10.1109/IPDPS.2016.61

We present a space and time efficient practical parallel algorithm for approximating the diameter of massive weighted undirected graphs on distributed-memory platforms supporting a MapReduce-like abstraction. The core of the algorithm is a randomized clustering strategy generating disjoint clusters of bounded weighted radius. Theoretically, our algorithm uses linear space and yields a polylogarithmic approximation guarantee; moreover, for the important class of graphs of bounded doubling dimension, it runs in a number of rounds asymptotically smaller than those required by the state-of-the-art Delta-stepping algorithm, which is the only practical linear-space competitor on distributed-memory platforms. We complement our theoretical findings with an extensive experimental analysis on large benchmark graphs, which demonstrates that our algorithm attains substantial improvements on a number of

key performance indicators with respect to the aforementioned competitor, while featuring a similar approximation ratio (a small constant less than 1.4, as opposed to the polylogarithmic theoretical bound).

We also point out some recent interesting applications of clustering for the efficient sequential approximation of diameter and node centralities in graphs of low doubling dimension.

4.22 Rademacher Averages: Theory and Practice

Matteo Riondato (Two Sigma Investments LP – New York, US)

Joint work of Matteo Riondato, Eli Upfal
Main reference M. Riondato, E. Upfal, "ABRA: Approximating Betweenness Centrality in Static and Dynamic Graphs with Rademacher Averages", in Proc. of the 22nd ACM SIGKDD Int'l Conf. on Knowledge Discovery and Data Mining (KDD 2016), pp. 1145–1154, ACM, 2016.
URL http://dx.doi.org/10.1145/2939672.2939770

Rademacher averages are fundamental concepts from statistical learning theory. They allow to study the maximum deviation of a family of sample averages from their expectations, using only information obtained from the sample itself. They were for long time considered to be only of theoretical interest, but in recent works we showed how to use them to speed up important tasks from data analysis, such as frequent pattern mining and centrality estimation from large graphs.

In this talk I first introduce the fundamental definitions and results involving Rademacher averages, including the celebrated symmetrization lemma and Massart's finite class lemma. Then I show how to use these results in practice to compute high-quality approximations of the betweenness centralities of all nodes in a graph by using progressive random sampling. This second part is based on joint work with Eli Upfal, appearing in the proceedings of ACM KDD'16.

4.23 Smoothed Analysis of Local Search for the Maximum-Cut Problem

Heiko Röglin (Universität Bonn, DE)

Joint work of Michael Etscheid, Heiko Röglin
Main reference M. Etscheid, H. Röglin, "Smoothed Analysis of Local Search for the Maximum-Cut Problem," ACM Transactions on Algorithms, 13(2):25, 2017.
URL http://dx.doi.org/10.1145/3011870

Even though local search heuristics are the method of choice in practice for many well-studied optimization problems, most of them behave poorly in the worst case. This is in particular the case for the Maximum-Cut Problem, for which local search can take an exponential number of steps to terminate and the problem of computing a local optimum is PLS-complete. To narrow the gap between theory and practice, we study local search for the Maximum-Cut Problem in the framework of smoothed analysis in which inputs are subject to a small amount of random noise. We show that the smoothed number of iterations is quasi-polynomial, i.e., it is bounded from above by a polynomial in $n^{\log n}$ and ϕ where n denotes the number of nodes and ϕ denotes the perturbation parameter. This shows that worst-case instances are fragile and it is a first step in explaining why they are rarely observed in practice.

4.24 On coalescence time in graphs – When is coalescing as fast as meeting?

Thomas Sauerwald (University of Cambridge, GB)

Joint work of Frederik Mallmann-Trenn, Varun Kanade, Thomas Sauerwald
Main reference V. Kanade, F. Mallmann-Trenn, T. Sauerwald, "On coalescence time in graphs – When is coalescing as fast as meeting?", arXiv:1611.02460 [cs.DM], 2016.
URL https://arxiv.org/abs/1611.02460

Coalescing random walks is a fundamental stochastic process, where a set of particles perform independent discrete-time random walks on an undirected graph. Whenever two or more particles meet at a given node, they merge and continue as a single random walk. The coalescence time is defined as the expected time until only one particle remains, starting from one particle at every node. Despite recent progress such as by Cooper et al. and Berenbrink et al., the coalescence time for graphs such as binary trees, d-dimensional tori, hypercubes and more generally, vertex-transitive graphs, remains unresolved.

We provide a powerful toolkit that results in tight bounds for various topologies including the aforementioned ones. The meeting time is defined as the worst-case expected time required for two random walks to arrive at the same node at the same time. As a general result, we establish that for graphs whose meeting time is only marginally larger than the mixing time (a factor of $\log^2 n$), the coalescence time of n random walks equals the meeting time up to constant factors. This upper bound is complemented by the construction of a graph family demonstrating that this result is the best possible up to constant factors. For almost-regular graphs, we bound the coalescence time by the hitting time, resolving the discrete-time variant of a conjecture by Aldous for this class of graphs. Finally, we prove that for any graph the coalescence time is bounded by $O(n^3)$. By duality, our results give bounds on the voter model.

4.25 Open Problem: k-means++ in data streams

Melanie Schmidt (Universität Bonn, DE)

The k-means++ algorithm is a simple yet powerful method to compute good solutions for the k-means problem. It beats the immensely popular heuristic for k-means, Lloyd's algorithm, in practice while providing a (moderate) approximation guarantee. However, it shares an issue with Lloyd's algorithm: it requires multiple passes over the data and does not scale too well. Modern clustering offers very strong theoretical results for k-means in the one pass setting and also practical algorithms that (heuristically) compute good solutions on the fly. However, it lacks an algorithm that is as elegant as the k-means++ algorithm.

4.26 Dependent Rounding

Aravind Srinivasan (University of Maryland – College Park, US)

Rounding algorithms that introduce carefully-chosen dependencies amongst the variables being rounded, have emerged as a useful technique in randomized (approximation) algorithms over the last twenty years or so. I will give a very brief overview of some aspects of this approach.

4.27 Greedy Algorithms for Stochastic Unrelated Machine Scheduling

Marc Uetz (University of Twente, NL)

Joint work of Varun Gupta, Ben Moseley, Marc Uetz, Qiaomin Xie
Main reference V. Gupta, B. Moseley, M. Uetz, Q. Xie, "Stochastic Online Scheduling on Unrelated Machines", in Proc. of the 19th Int'l Conf. on Integer Programming and Combinatorial Optimization (IPCO 2017), LNCS, Vol. 10328, pp. 228–240, Springer, 2017.
URL http://dx.doi.org/10.1007/978-3-319-59250-3_19

We derive the first performance guarantees for an online algorithm that schedules stochastic, nonpreemptive jobs on unrelated machines to minimize the expectation of the total weighted completion time. The performance guarantees depend linearly on the squared coefficient of variation of the underlying random variables, and therefore –asymptotically– match previously known results. Prior work on unrelated machine scheduling was either restricted to the offline case, and/or required sophisticated linear or convex programming relaxations for the assignment of jobs to machines. In contrast, our algorithm is purely combinatorial, and therefore it works for the online setting, too. As to the techniques applied, we obtain our results by dual fitting. In that sense, we also demonstrate how this technique be put to work for stochastic and nonpreemptive scheduling problems.

4.28 Reconstructing Hidden Permutations Using the Average-Precision (AP) Correlation Statistic

Eli Upfal (Brown University – Providence, US)

Joint work of Lorenzo De Stefani, Alessandro Epasto, Eli Upfal, Fabio Vandin
Main reference L. De Stefani, A. Epasto, E. Upfal, F. Vandin, "Reconstructing Hidden Permutations Using the Average-Precision (AP) Correlation Statistic", in Proc. of the 31st AAAI Conf. on Artificial Intelligence (AAAI 2017), pp. 1526–1532, 2016.
URL http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12046

We study the problem of learning probabilistic models for permutations, where the order between highly ranked items in the observed permutations is more reliable (i.e., consistent in different rankings) than the order between lower ranked items, a typical phenomena observed in many applications such as web search results and product ranking. We introduce and study a variant of the Mallows model where the distribution is a function of the widely used Average-Precision (AP) Correlation statistic, instead of the standard Kendall's tau distance.

We present a generative model for constructing samples from this distribution and prove useful properties of that distribution. Using these properties we develop an efficient algorithm that provably computes an asymptotically unbiased estimate of the center permutation, and a faster algorithm that learns with high probability the hidden central permutation for a wide range of the parameters of the model. We complement our theoretical analysis with extensive experiments showing that unsupervised methods based on our model can precisely identify ground-truth clusters of rankings in real-world data. In particular, when compared to the Kendall's tau based methods, our methods are less affected by noise in low-rank items.

4.29 Randomized Shared Memory Algorithms

Philipp Woelfel (University of Calgary, CA)

Joint work of George Giakkoupis, Philipp Woelfel
Main reference G. Giakkoupis, P. Woelfel, "Randomized Mutual Exclusion with Constant Amortized RMR Complexity on the DSM", in Proc. of the 55th Annual Symposium on Foundations of Computer Science (FOCS 2014), pp. 504–513, IEEE, 2014.
URL http://dx.doi.org/10.1109/FOCS.2014.60
Main reference G. Giakkoupis, P. Woelfel, "On the time and space complexity of randomized test-and-set", in Proc. of the 2012 ACM symposium on Principles of distributed computing (PODC 2012), pp. 19–28, ACM, 2012.
URL http://dx.doi.org/10.1145/2332432.2332436

Randomization has been recognized for decades as a powerful tool in sequential algorithm design. But for concurrent algorithmic problems, probabilistic choices by processes have been mainly used to solve otherwise unsolvable problems (most famously the consensus problem). Only in recent years, randomization has received attention to improve the efficiency of concurrent algorithms.

In this talk, I will give an overview of the asynchronous shared memory model, and how randomization can be employed to find solutions to some fundamental problems, which are more efficient than the best known deterministic algorithms. I will give two main examples of recently devised randomized algorithms:

- A randomized leader election algorithm that has $O(\log \log n)$ expected step complexity; and
- A randomized mutual exclusion algorithm that has amortized constant expected RMR complexity.

Participants

- Dimitris Achlioptas
University of California – Santa Cruz, US
- Aris Anagnostopoulos
Sapienza University of Rome, IT
- Luca Becchetti
Sapienza University of Rome, IT
- Petra Berenbrink
Universität Hamburg, DE
- Amin Coja-Oghlan
Goethe-Universität – Frankfurt am Main, DE
- Artur Czumaj
University of Warwick – Coventry, GB
- Daniel Dadush
CWI – Amsterdam, NL
- Martin Dietzfelbinger
TU Ilmenau, DE
- Benjamin Doerr
Ecole Polytechnique – Palaiseau, FR
- Devdatt Dubhashi
Chalmers UT – Göteborg, SE
- Matthias Englert
University of Warwick – Coventry, GB
- Leslie Ann Goldberg
University of Oxford, GB
- Paul W. Goldberg
University of Oxford, GB
- Ross Kang
Radboud University Nijmegen, NL
- Thomas Kesselheim
TU Dortmund, DE
- Stefan Klootwijk
University of Twente, NL
- Robert Krauthgamer
Weizmann Institute – Rehovot, IL
- Marvin Künnemann
MPI für Informatik – Saarbrücken, DE
- Stefano Leonardi
Sapienza University of Rome, IT
- Frederik Mallmann-Trenn
ENS – Paris, FR
- Bodo Manthey
University of Twente, NL
- Claire Mathieu
ENS – Paris, FR
- Nicole Megow
Universität Bremen, DE
- Morteza Monemizadeh
Charles University – Prague, CZ
- Seffi Naor
Technion – Haifa, IL
- Alessandro Panconesi
Sapienza University of Rome, IT
- Andrea Pietracaprina
University of Padova, IT
- Kirk Pruhs
University of Pittsburgh, US
- Geppino Pucci
University of Padova, IT
- Matteo Riondato
Two Sigma Investments LP – New York, US
- Heiko Röglin
Universität Bonn, DE
- Thomas Sauerwald
University of Cambridge, GB
- Melanie Schmidt
Universität Bonn, DE
- Christian Sohler
TU Dortmund, DE
- Aravind Srinivasan
University of Maryland – College Park, US
- Marc Uetz
University of Twente, NL
- Eli Upfal
Brown University – Providence, US
- Tjark Vredeveld
Maastricht University, NL
- Philipp Woelfel
University of Calgary, CA