Final Program

Association for Computing Machinery

ACM-SIAM Symposium on Discrete Algorithms

January 6-9, 2019
Westin San Diego
San Diego, California, USA

ALENEX19

Meeting on
Algorithm Engineering & Experiments

January 7-8, 2019
Westin San Diego
San Diego, California, USA

ANALCO19

Meeting on Analytic Algorithmics and Combinatorics

January 6, 2019
Westin San Diego
San Diego, California, USA

The Symposium on Simplicity in Algorithms (SOSA19) will take place on January 8 and 9.

SODA is jointly sponsored by the ACM Special Interest Group on Algorithms and Computation Theory, and the SIAM Activity Group on Discrete Mathematics

The SIAM Activity Group on Discrete Mathematics focuses on combinatorics, graph theory, cryptography, discrete optimization, mathematical programming, coding theory, information theory, game theory, and theoretical computer science, including algorithms, complexity, circuit design, robotics, and parallel processing. We provide an opportunity to unify pure discrete mathematics and areas of applied research such as computer science, operations research, combinatorics, and the social sciences.

Society for Industrial and Applied Mathematics

3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
Telephone: +1-215-382-9800 Fax: +1-215-386-7999
Conference E-mail: meetings@siam.org
Conference Web: *siam.org/meetings*
Membership and Customer Service:
(800) 447-7426 (USA & Canada) or +1-215-382-9800 (worldwide)
https://www.siam.org/conferences/CM/Main/soda19
https://www.siam.org/conferences/CM/Main/alenex19
https://www.siam.org/conferences/CM/Main/analco19
https://simplicityinalgorithms.com/

Table of Contents

SODA COMMITTEES
Program Committee Chair

Timothy M. Chan
University of Illinois at Urbana-Champaign, U.S.

Program Committee

Peyman Afshani
Aarhus University, Denmark
Shipra Agrawal
Columbia University, U.S.
Aditya Bhaskara
University of Utah, U.S.
Yang Cai
McGill University, Canada
Amit Chakrabarti
Dartmouth College, U.S.
Deeparnab Chakrabarty
Dartmouth College, U.S.
Erin Wolf Chambers
Saint Louis University, U.S.
Siu On Chan
Chinese University of Hong Kong, Hong Kong
Daniel Dadush
CWI, Netherlands
Holger Dell
Saarland University, Germany
Zdenek Dvorak
Charles University, Czech Republic
Friedrich Eisenbrand
EPFL, Switzerland
Michael Elkin
Ben-Gurion University of the Negev, Israel
Jane Gao
Monash University, Australia
Sariel Har-Peled
University of Illinois at Urbana-Champaign, U.S.
Hamed Hatami
McGill University, Canada
Thomas Hayes
University of New Mexico, U.S.
Martin Hoefer
Goethe University Frankfurt, Germany
Thore Husfeldt
Lund University, Sweden and IT University of Copenhagen, Denmark
Klaus Jansen
Christian-Albrechts-Universität zu Kiel, Germany
T.S. Jayram
IBM Almaden Research, U.S.
Daniel Kane
University of California, San Diego, U.S.
Jonathan Kelner
Massachusetts Institute of Technology, U.S.
Tsvi Kopelowitz
Bar Ilan University, Israel
Pravesh K. Kothari
Princeton University and the Institute for Advanced Study, U.S.
Fabian Kuhn
University of Freiburg, Germany
Daniel Lokshtanov
University of Bergen, Norway
Brendan Lucier
Microsoft Research, U.S.
Bojan Mohar
Simon Fraser University, Canada
Viswanath Nagarajan
University of Michigan, U.S.
Alantha Newman
CNRS, Grenoble, France
Ilan Newman
University of Haifa, Israel
Vijaya Ramachandran
University of Texas at Austin, U.S.
Ilya Razenshteyn
Microsoft Research, U.S.
Liam Roditty
Bar Ilan University, Israel
Atri Rudra
University at Buffalo, SUNY, U.S.
Laura Sanità
University of Waterloo, Canada
László A. Végh
London School of Economics, United Kingdom
Oren Weimann
University of Haifa, Israel
Udi Wieder
VMware Research, U.S.

Steering Committee

Pavol Hell
Simon Fraser University, Canada
Daniel Král
University of Warwick, United Kingdom
Dana Randall
Georgia Institute of Technology, U.S.
Cliff Stein
Columbia University, U.S. (chair)
Shang-Hua Teng
University of Southern California, U.S.

ALENEX COMMITTEES
Program Committee Co-Chairs

Stephen Kobourov
University of Arizona, U.S.
Henning Meyerhenke
Humboldt-Universität zu Berlin, Germany

Program Committee

Aaron Archer
Google, U.S.
Maike Buchin
Ruhr Universität Bochum, Germany
Aydin Buluc
Lawrence Berkeley National Laboratory, U.S.
Markus Chimani
Osnabrück University, Germany
Pierluigi Crescenzi
Università degli Studi di Firenze, Italy
Yifan Hu
Yahoo!, U.S.
Yoichi Iwata
National Institute of Informatics, Japan
Juha Kärkkäinen
University of Helsinki, Finland
Ken-Ichi Kawarabayashi
National Institute of Informatics, Japan
Stefan Kratsch
Humboldt-Universität zu Berlin, Germany
Fredrik Manne
University of Bergen, Norway
Nicole Megow
Universität Bremen, Germany
Ulrich Meyer
Goethe University Frankfurt, Germany
Matthias Müller-Hannemann
Martin-Luther-University Halle-Wittenberg, Germany
Matthias Petri
The University of Melbourne, Australia
Sergey Pupyrev
Facebook, U.S.
Barna Saha
University of Massachusetts, Amherst, U.S.
Renato Werneck
Amazon, U.S.

Steering Committee

Andrew V. Goldberg
Amazon.com, U.S. (Chair)
Michael Goodrich
University of California, Irvine, U.S.

Giuseppe F. Italiano
University of Rome "Tor Vergata", Italy
Rasmus Pagh
IT University of Copenhagen, Denmark
Vijaya Ramachandran
University of Texas, Austin, U.S.
Cliff Stein
Columbia University, U.S.
Dorothea Wagner
Karlsruhe Institute of Technology, Germany

ANALCO COMMITTEES

Program Committee Co-Chairs

Marni Mishna
Simon Fraser University, Canada
J. Ian Munro
University of Waterloo, Canada

Program Committee

Martin Dietzfelbinger
Technische Universität Ilmenau, Germany
Cecilia Holmgren
Uppsala University, Sweden
Mihyun Kang
Technische Universität Graz, Austria
Yusuke Kobayashi
Kyoto University, Japan
Jérémie Lumbroso
Princeton University, U.S.
Hosam Mahmoud
George Washington University, U.S.
Daniel Panario
Carleton University, Canada
Dominique Poulalhon
Université Paris Diderot, France
Sebastian Wild
University of Waterloo, Canada

Steering Committee

Michael Drmota
Technische Universität Wien, Austria
James Allen Fill (Jim Fill)
Johns Hopkins University, U.S.
H.K. Hwang
Institute of Statistical Science, Academia Sinica, Taiwan
Conrado Martínez
Universitat Politècnica de Catalunya, Spain
Markus Nebel
Universität Bielefeld, Germany
Robert Sedgewick
Princeton University, U.S.
Wojciech Szpankowski
Purdue University, U.S.
Mark Daniel Ward
Purdue University, U.S.

SODA Themes

Aspects of combinatorics and discrete mathematics, such as:
Combinatorial structures
Discrete optimization
Discrete probability
Finite metric spaces
Graph theory
Mathematical programming
Random structures
Topological problems

Core topics in discrete algorithms, such as:
Algorithm analysis
Data structures
Experimental algorithmics

Algorithmic aspects of other areas of computer science, such as:
Combinatorial scientific computing
Communication networks and the Internet
Computational geometry and topology
Computer graphics and computer vision
Computer systems
Cryptography and security
Databases and information retrieval
Data compression
Data privacy
Distributed and parallel computing
Game theory and mechanism design
Machine learning
Quantum computing

SIAM Registration Desk

The registration desk is on the 2nd Floor. It is open during the following hours:

Saturday, January 5
5:00 p.m. – 8:00 p.m.

Sunday, January 6
8:00 a.m. – 5:00 p.m.

Monday, January 7
8:00 a.m. – 5:00 p.m.

Tuesday, January 8
8:00 a.m. – 5:00 p.m.

Wednesday, January 9
8:00 a.m. – 5:00 p.m.

Hotel Address

The Westin San Diego
400 West Broadway
San Diego, CA, 92101, U.S.

Hotel Telephone Number

To reach an attendee or leave a message, call +1-619-239-4500. If the attendee is a hotel guest, the hotel operator can connect you with the attendee's room.

Hotel Check-in and Check-out Times

Check-in time is 3:00 p.m.
Check-out time is 12:00 p.m.

Child Care

California DMC recommends Destinations Sitters (*https://www.destinationsitters.com/*) for attendees interested in child care services. Attendees are responsible for making their own child care arrangements.

Corporate/Institutional Members

SIAM corporate members provide their employees with knowledge about, access to, and contacts in the applied mathematics and computational sciences community through their membership benefits. Corporate membership is more than just a bundle of tangible products and services; it is an expression of support for SIAM and its programs. SIAM is pleased to acknowledge its corporate members and sponsors. In recognition of their support, non-member attendees who are employed by the following organizations are entitled to the SIAM member registration rate.

Corporate/Institutional Members

The Aerospace Corporation

Air Force Office of Scientific Research

Amazon

Argonne National Laboratory

Bechtel Marine Propulsion Laboratory

The Boeing Company

CEA/DAM

Cirrus Logic

Department of National Defence (DND/CSEC)

DSTO- Defence Science and Technology Organisation, Edinburgh

Exxon Mobil

IDA Center for Communications Research, La Jolla

IDA Institute for Defense Analyses, Princeton

IDA Institute for Defense Analyses, Bowie, Maryland

Lawrence Berkeley National Laboratory

Lawrence Livermore National Labs

Lockheed Martin Maritime Systems & Sensors

Los Alamos National Laboratory

Max-Planck-Institute

Mentor Graphics

National Institute of Standards and Technology (NIST)

National Security Agency

Oak Ridge National Laboratory

Sandia National Laboratories

Schlumberger

Simons Foundation

United States Department of Energy

U.S. Army Corps of Engineers, Engineer Research and Development Center

List current November 2018.

Join SIAM and save!

Leading the applied mathematics community . . .

SIAM members save up to $140 on full registration for SODA and its associated meetings. Join your peers in supporting the premier professional society for applied mathematicians and computational scientists. SIAM members receive subscriptions to *SIAM Review, SIAM News* and *SIAM Unwrapped*, and enjoy substantial discounts on SIAM books, journal subscriptions, and conference registrations.

If you are not a SIAM member and did not join SIAM before you registered, you can apply the difference of $140 between what you paid as a non-member, and what a member paid towards a SIAM membership. Contact SIAM Customer Service for details or join at the conference registration desk.

If you are a SIAM member, it only costs $15 to join the SIAM Activity Group on Discrete Mathematics.

Students who paid the Student Non Member Rate will be automatically enrolled as SIAM Student Members. Please go to *my.siam.org* to update your education and contact information in your profile. If you attend a SIAM Academic Member Institution or are part of a SIAM Student Chapter you will be able to renew next year for free.

Join onsite at the registration desk, go to *https://www.siam.org/Membership/Join-SIAM* to join online or download an application form, or contact SIAM Customer Service:
Telephone: +1-215-382-9800 (worldwide); or 800-447-7426 (U.S. and Canada only)
Fax: +1-215-386-7999
Email: membership@siam.org
Postal mail: Society for Industrial and Applied Mathematics, 3600 Market Street, 6th floor, Philadelphia, PA 19104-2688 U.S.

Standard Audio-Visual Set-Up in Meeting Rooms

SIAM does not provide computers for any speaker. When giving an electronic presentation, speakers must provide their own computers. SIAM is not responsible for the safety and security of speakers' computers.

A data (LCD) projector and screen will be provided in all technical session meeting rooms. The data projectors support both VGA and HDMI connections. Presenters requiring an alternate connection must provide their own adaptor.

Internet Access

Attendees booked within the SIAM room block will receive complimentary wireless internet access in their guest rooms and the public areas of the hotel. All conference attendees will have complimentary wireless internet access in the meeting space of the hotel.

SIAM will provide a limited number of email stations for attendees during registration hours.

Registration Fee Includes

- Admission to all technical sessions (SODA19, ALENEX19, ANALCO19, SOSA19)
- ANALCO/ALENEX and SOSA Business meetings
- Coffee breaks daily
- Continental Breakfast daily
- Luncheon on Sunday, January 6
- Proceedings (SODA, ALENEX, and ANALCO posted online)
- Room set-ups and audio/visual equipment
- SODA Business Meeting
- Welcome Reception

Job Postings

Please check with the SIAM registration desk regarding the availability of job postings or visit *https://jobs.siam.org/*.

Table Top Display

Cambridge University Press

Conference Sponsors

Google IBM Research

Julia computing

2019 Conference Bag Sponsor

Microsoft SIGACT

Name Badges

A space for emergency contact information is provided on the back of your name badge. Help us help you in the event of an emergency!

Comments?

Comments about SIAM meetings are encouraged! Please send to:
Cynthia Phillips, SIAM Vice President for Programs (*vpp@siam.org*).

Get-togethers

Welcome Reception
Saturday, January 5
6:00 p.m. – 8:00 p.m.
Poolside and Ivory - 3rd Floor

ALENEX and ANALCO Business Meeting
Sunday, January 6
6:45 p.m. – 7:45 p.m.
Diamond 2 - 2nd Floor

SODA Business Meeting and Awards Presentation
Monday, January 7
6:45 p.m. – 7:45 p.m.
Emerald Ballroom - 2nd Floor
Complimentary beer and wine will be served.

SOSA Business Meeting
Tuesday, January 8
6:45 p.m. – 7:45 p.m.
Diamond 2 - 2nd Floor

Statement on Inclusiveness

As a professional society, SIAM is committed to providing an inclusive climate that encourages the open expression and exchange of ideas, that is free from all forms of discrimination, harassment, and retaliation, and that is welcoming and comfortable to all members and to those who participate in its activities. In pursuit of that commitment, SIAM is dedicated to the philosophy of equality of opportunity and treatment for all participants regardless of gender, gender identity or expression, sexual orientation, race, color, national or ethnic origin, religion or religious belief, age, marital status, disabilities, veteran status, field of expertise, or any other reason not related to scientific merit. This philosophy extends from SIAM conferences, to its publications, and to its governing structures and bodies. We expect all members of SIAM and participants in SIAM activities to work towards this commitment.is welcoming and comfortable to all members and to those who participate in its activities. In pursuit of that commitment, SIAM is dedicated to the philosophy of equality of opportunity and treatment for all participants regardless of gender, gender identity or expression, sexual orientation, race, color, national or ethnic origin, religion or religious belief, age, marital status, disabilities, veteran status, field of expertise, or any other reason not related to scientific merit. This philosophy extends from SIAM conferences, to its publications, and to its governing structures and bodies. We expect all members of SIAM and participants in SIAM activities to work towards this commitment.

Please Note

SIAM is not responsible for the safety and security of attendees' computers. Do not leave your laptop computers unattended. Please remember to turn off your cell phones, pagers, etc. during sessions.

Recording of Presentations

Audio and video recording of presentations at SIAM meetings is prohibited without the written permission of the presenter and SIAM.

Social Media

SIAM is promoting the use of social media, such as Facebook and Twitter, in order to enhance scientific discussion at its meetings and enable attendees to connect with each other prior to, during and after conferences. If you are tweeting about a conference, please use the designated hashtag to enable other attendees to keep up with the Twitter conversation and to allow better archiving of our conference discussions. The hashtags for these meetings are #SIAMDA19, #ANALCO19 and #ALENEX19.

SIAM's Twitter handle is @TheSIAMNews.

Changes to the Printed Program

The printed program was current at the time of printing, however, please review the online program schedule (*http://meetings.siam.org/program.cfm?CONFCODE=da19*) for up-to-date information.

Gene Golub g2s3 2019 SIAM Summer School

June 17-28, 2019
Aussois, France

HIGH PERFORMANCE DATA ANALYTICS

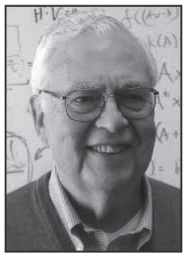

The tenth Gene Golub SIAM Summer School will take place in France, at the Paul Langevin conference center in Aussois, in the French Alps.

The focus of the school will be on large-scale data analytics, which lies at the intersections of data analytics algorithms and high performance computing. Students will be introduced to problems in data analytics arising from both the machine learning and the scientific computing communities. The school will include perspectives from industry, such as Amazon, Google, and IBM, as well as from academic instructors.

Students will be exposed to "end-to-end" multidisciplinary topics, which span several traditionally disparate areas. The series of lectures will develop background on methods and algorithms for data analytics, approximation algorithms to deal with large volumes of data, languages and tools for implementing those algorithms on large scale computers, and data-driven applications from scientific computing and machine learning.

The summer school is being organized by Laura Grigori (Inria and Sorbonne University), Matthew Knepley (University at Buffalo) Olaf Schenk (Università della Svizzera Italiana), and Rich Vuduc (Georgia Institute of Technology).

The intended audience is intermediate graduate students (students with a Master's degree, 2nd-3rd year Ph.D. students without an MS, or equivalent). Applicants selected to participate pay no registration fee. Funding for local accommodations and meal expenses will be available for all participants.

Application deadline: February 1, 2019

As information becomes available on how to apply, it will be posted at:

http://www.siam.org/students/g2s3/

Sponsored by SIAM through an endowment from the estate of Gene Golub.

For more information about prior summer schools and Professor Gene Golub go to ***http://www.siam.org/students/g2s3/***

siam®

Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor • Philadelphia, PA 19104-2688 USA • +1-215-382-9800
siam@siam.org • www.siam.org

SIAM PRESENTS IS AN AUDIO-VISUAL ARCHIVE

COMPRISED OF MORE THAN 2,000 PRESENTATIONS POSTED IN OVER 40 SEARCHABLE TOPICS, INCLUDING:

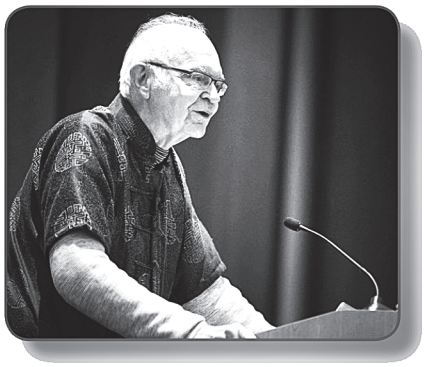

- **algebraic geometry**
- **atmospheric and oceanographic science**
- **computational science**
- **data mining**
- **geophysical science**
- **optimization**
- **uncertainty quantification and more...**

The collection, *Featured Lectures from our Archives*, includes audio and slides from more than 30 conferences since 2008, including talks by invited and prize speakers, select minisymposia, and minitutorials. Presentations from SIAM meetings are being added throughout the year.

In addition you can view short video clips of speaker interviews from sessions at Annual Meetings starting in 2010.

Plans for adding more content are on the horizon. Keep an eye out!

The audio, slide, and video presentations are part of SIAM's outreach activities to increase the public's awareness of mathematics and computational science in the real world, and to bring attention to exciting and valuable work being done in the field. Funding from SIAM, the National Science Foundation, and the Department of Energy was used to partially support this project.

New presentations are posted every few months as the program expands with sessions from additional SIAM meetings. Users can search for presentations by category, speaker name, and/or key words.

www.siam.org/meetings/presents.php

siam®
SOCIETY FOR INDUSTRIAL AND APPLIED MATHEMATICS · 3600 MARKET STREET, 6TH FLOOR · PHILADELPHIA, PA 19104-2688 USA
PHONE: +1-215-382-9800 · FAX +1-215-386-7999 · SERVICE@SIAM.ORG · WWW.SIAM.ORG

Best Paper and Best Student Paper Awards

Best Paper Award

Sublinear Algorithms for (Δ + 1) Vertex Coloring

Sepehr Assadi, Yu Chen, and Sanjeev Khanna, ***University of Pennsylvania, U.S.***

This paper will be presented on Monday, January 7,
in session CP14 SODA Session 4A.

See page 18 for session details.

Best Student Paper Award

Optimal Las Vegas Approximate Near Neighbors in ℓ_p

Alexander Wei, *Harvard University, U.S.*

This paper will be presented on Tuesday, January 8,
in session CP30 SODA Session 8A.

See page 25 for session details.

Invited Plenary Speakers

All Invited Plenary Presentations will take place in Emerald Ballroom - 2nd Floor

Sunday, January 6

11:30 a.m. - 12:30 p.m

IP1 Towards Analysis of Information Content in Dynamic Networks
Wojciech Szpankowski, *Purdue University, U.S.*

Monday, January 7

11:30 a.m. - 12:30 p.m.

IP2 Recent Advances on the Complexity of Parameterized Counting Problems
Dániel Marx, *Hungarian Academy of Sciences, Hungary*

Tuesday, January 8

11:30 a.m. - 12:30 p.m.

IP3 Inherent Trade-Offs in Algorithmic Fairness
Jon M. Kleinberg, *Cornell University, U.S.*

Wednesday, January 9

11:30 a.m. - 12:30 p.m.

IP4 (Hardness of) Approximation and Expansion
Irit Dinur, *Weizmann Institute of Science, Israel*

Program Schedule

siam®

acm Association for Computing Machinery

ACM-SIAM Symposium on Discrete Algorithms

January 6-9, 2019
Westin San Diego
San Diego, California, USA

ALENEX19

Meeting on
Algorithm Engineering & Experiments

January 7-8, 2019
Westin San Diego
San Diego, California, USA

ANALCO19

Meeting on Analytic Algorithmics and Combinatorics

January 6, 2019
Westin San Diego
San Diego, California, USA

Symposium on Simplicity in Algorithms (SOSA19)

Saturday, January 5

Registration

5:00 PM-8:00 PM

Room: Registration Desk - 2rd Floor

Welcome Reception

6:00 PM-8:00 PM

Room: Poolside and Ivory - 3rd Floor

Sunday, January 6

Registration

8:00 AM-5:00 PM

Room: Registration Desk - 2rd Floor

Continental Breakfast

8:30 AM-9:00 AM

Room: Ballroom Foyer - 2nd Floor

Sunday, January 6

CP1

ANALCO Session 1

9:00 AM-11:05 AM

Room: Diamond 2 - 2nd Floor

Chair: Sebastian Wild, University of Waterloo, Canada

9:00-9:20 Combinatorics of Nondeterministic Walks of the Dyck and Motzkin Type

Elie de Panafieu, Nokia Bell Labs France, France; Mohamed Lamine Lamali and Michael Wallner, LaBRI - Université de Bordeaux, France

9:25-9:45 Ranked Schröder Trees

Antoine Genitrini, Sorbonne Universités, France; Olivier Bodini, Université Paris 13, France; Mehdi Naima, Universite Paris-Nord, France

9:50-10:10 Esthetic Numbers and Lifting Restrictions on the Analysis of Summatory Functions of Regular Sequences

Clemens Heuberger and Daniel Krenn, Alpen-Adria-Universität Klagenfurt, Austria

10:15-10:35 Reducing Simply Generated Trees by Iterative Leaf Cutting

Benjamin Hackl and Clemens Heuberger, Alpen-Adria-Universität Klagenfurt, Austria; Stephan Wagner, Stellenbosch University, South Africa

10:40-11:00 Protection Number of Recursive Trees

Zbigniew Golebiewski and *Mateusz Klimczak*, Wroclaw University of Science and Technology, Poland

Sunday, January 6

CP2

SODA Session 1A

9:00 AM-11:05 AM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

9:00-9:20 Fine-grained Complexity Meets IP = PSPACE
Lijie Chen and Shafi Goldwasser, Massachusetts Institute of Technology, U.S.; *Kaifeng Lyu*, Tsinghua University, China; Guy Rothblum, Microsoft Research, U.S.; Aviad Rubinstein, Harvard University, U.S.

9:25-9:45 An Equivalence Class for Orthogonal Vectors
Lijie Chen and Ryan Williams, Massachusetts Institute of Technology, U.S.

9:50-10:10 Seth-Based Lower Bounds for Subset Sum and Bicriteria Path
Amir Abboud, IBM Almaden Research Center, U.S.; *Karl Bringmann*, Max Planck Institute for Informatics, Germany; Danny Hermelin and Dvir Shabtay, Ben-Gurion University, Israel

10:15-10:35 Fast Modular Subset Sum Using Linear Sketching
Kyriakos Axiotis, Massachusetts Institute of Technology, U.S.; Arturs Backurs, Toyota Technological Institute at Chicago, U.S.; Ce Jin, Tsinghua University, China; Christos Tzamos, University of Wisconsin, Madison, U.S.; Hongxun Wu, Tsinghua University, China

10:40-11:00 A Subquadratic Approximation Scheme for Partition
Marcin Mucha, *Karol Wegrzycki*, and Michal Wlodarczyk, University of Warsaw, Poland

Sunday, January 6

CP3

SODA Session 1B

9:00 AM-11:05 AM

Room: Topaz - 2nd Floor

Chair: To Be Determined

9:00-9:20 Metrical Task Systems on Trees Via Mirror Descent and Unfair Gluing
Sebastien Bubeck, Microsoft Research, U.S.; Michael B. Cohen, Massachusetts Institute of Technology, U.S.; James R. Lee and Yin Tat Lee, University of Washington, U.S.

9:25-9:45 *K*-Servers with a Smile: Online Algorithms via Projections
Marco Molinaro, Pontifical Catholic University of Rio de Janeiro, Brazil; Niv Buchbinder, Tel Aviv University, Israel; Anupam Gupta, Carnegie Mellon University, U.S.; Joseph Naor, Technion Israel Institute of Technology, Israel

9:50-10:10 A Nearly-Linear Bound for Chasing Nested Convex Bodies
C.J. Argue, Carnegie Mellon University, U.S.; Sebastien Bubeck, Microsoft Research, U.S.; Michael B. Cohen, Massachusetts Institute of Technology, U.S.; Anupam Gupta, Carnegie Mellon University, U.S.; Yin Tat Lee, University of Washington, U.S.

10:15-10:35 A Ø-Competitive Algorithm for Scheduling Packets with Deadlines
Pavel Vesely, University of Warwick, United Kingdom; Marek Chrobak, University of California, Riverside, U.S.; Lukasz Jez, Tel Aviv University, Israel and University of Wroclaw, Poland; Jiri Sgall, Charles University, Czech Republic

10:40-11:00 Elastic Caching
Debmalya Panigrahi, Duke University, U.S.; Anupam Gupta, Carnegie Mellon University, U.S.; Ravishankar Krishnaswamy, Microsoft Research, India; Amit Kumar, Indian Institute of Technology, Delhi, India

Sunday, January 6

CP4

SODA Session 1C

9:00 AM-11:05 AM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

9:00-9:20 Dynamic Double Auctions: Towards First Best
Santiago Balseiro, Columbia University, U.S.; Vahab Mirrokni, Google, Inc., U.S.; Renato Paes Leme and *Song Zuo*, Google Research, U.S.

9:25-9:45 Multi-unit Supply-monotone Auctions with Bayesian Valuations
Yuan Deng and Debmalya Panigrahi, Duke University, U.S.

9:50-10:10 Correlation-robust Analysis of Single Item Auction
Xiaohui Bei, Nanyang Technological University, Singapore; Nick Gravin and Pinyan Lu, Shanghai University of Finance and Economics, China; Zhihao Gavin Tang, University of Hong Kong, Hong Kong

10:15-10:35 Tight Revenue Gaps among Simple Mechanisms
Yaonan Jin, Hong Kong University of Science and Technology, Hong Kong; Pinyan Lu, Shanghai University of Finance and Economics, China; Zhihao Gavin Tang, University of Hong Kong, Hong Kong; Tao Xiao, Shanghai Jiao Tong University, China

10:40-11:00 Assignment Mechanisms under Distributional Constraints
Itai Ashlagi, Amin Saberi, and *Ali Shameli*, Stanford University, U.S.

Coffee Break

11:05 AM-11:30 AM

Room: Ballroom Foyer - 2nd Floor

Sunday, January 6

IP1

Towards Analysis of Information Content in Dynamic Networks

11:30 AM-12:30 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

Shannon information theory has served as a bedrock for advances in communication and storage systems over the past six decades. However, this theory does not handle well higher order structures (e.g., graphs, geometric structures), temporal aspects (e.g., real-time considerations), or semantics, which are essential aspects of data and information that underlie a broad class of current and emerging data science applications. In this talk, we present some recent results on structural and temporal information in dynamic networks/graphs, in which nodes and edges are added and removed over time. We focus on two related problems: (i) compression of structures -- for a given graph model, we exhibit an efficient algorithm for invertibly mapping network structures (i.e., graph isomorphism types) to bit strings of minimum expected length, and (ii) node arrival order inference -- for a dynamic graph model, we determine the extent to which the order of node arrivals can be inferred from a snapshot of the graph structure. For both problems, we apply analytic combinatorics, probabilistic, and information-theoretic methods to find statistical limits and efficient algorithms for achieving those limits.

Wojciech Szpankowski
Purdue University, U.S.

Luncheon **Ticketed Event**

12:30 PM-2:00 PM

Room:Crystal Ballroom - 2nd Floor

Please visit the SIAM Registration Desk if you did not receive a luncheon ticket.

Sunday, January 6

CP5

ANALCO Session 2

2:00 PM-3:40 PM

Room: Diamond 2 - 2nd Floor

Chair: Hosam Mahmoud, George Washington University, U.S.

2:00-2:20 Sesquickselect: One and a Half Pivots for Cache-Efficient Selection
Sebastian Wild, University of Waterloo, Canada; Conrado Martínez, Universidad Politecnica de Catalunya, Spain; Markus Nebel, Universität Bielefeld, Germany

2:25-2:45 Moments of Select Sets
Simon H. Langowski and Mark Daniel Ward, Purdue University, U.S.

2:50-3:10 Median-of-K Jumplists and Dangling-Min BSTs
Sebastian Wild, University of Waterloo, Canada; *Markus Nebel*, Universität Bielefeld, Germany; Elisabeth Neumann, Technische Universität Braunschweig, Germany

3:15-3:35 QuickSort: Improved Right-Tail Asymptotics for the Limiting Distribution, and Large Deviations
James Allen Fill and Wei-Chun Hung, Johns Hopkins University, U.S.

Sunday, January 6

CP6

SODA Session 2A

2:00 PM-4:05 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

2:00-2:20 Deterministic (1/2 + ε)-Approximation for Submodular Maximization over a Matroid
Niv Buchbinder, Tel Aviv University, Israel; Moran Feldman and *Mohit Garg*, Open University of Israel, Israel

2:25-2:45 Submodular Maximization with Nearly Optimal Approximation, Adaptivity and Query Complexity
Matthew Fahrbach, Georgia Institute of Technology, U.S.; Vahab Mirrokni and Morteza Zadimoghaddam, Google, Inc., U.S.

2:25-2:45 Submodular Maximization with Nearly-optimal Approximation and Adaptivity in Nearly-linear Time
Alina Ene, Boston University, U.S.; Huy Nguyen, Northeastern University, U.S.

2:25-2:45 An Exponential Speedup in Parallel Running Time for Submodular Maximization Without Loss in Approximation
Eric Balkanski, Aviad Rubinstein, and Yaron Singer, Harvard University, U.S.

2:50-3:10 Submodular Function Maximization in Parallel via the Multilinear Extension
Chandra Chekuri and *Kent Quanrud*, University of Illinois at Urbana-Champaign, U.S.

3:15-3:35 Stochastic Submodular Cover with Limited Adaptivity
Arpit Agarwal, Sepehr Assadi, and Sanjeev Khanna, University of Pennsylvania, U.S.

3:40-4:00 Stochastic ℓp Load Balancing and Moment Problems via the *L*-Function Method
Marco Molinaro, Pontifical Catholic University of Rio de Janeiro, Brazil

Sunday, January 6

CP7

SODA Session 2B

2:00 PM-4:05 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

2:00-2:20 Approximate Nearest Neighbor Searching with Non-euclidean and Weighted Distances

Ahmed Abdelkader, University of Maryland, U.S.; Sunil Arya, Hong Kong University of Science and Technology, Hong Kong; Guilherme D. da Fonseca, Université Clermont Auvergne, France; David M. Mount, University of Maryland, U.S.

2:25-2:45 Colored Range Closest-Pair Problem under General Distance Functions

Jie Xue, University of Minnesota, Twin Cities, U.S.

2:50-3:10 Optimal Algorithm for Geodesic Nearest-point Voronoi Diagrams in Simple Polygons

Eunjin Oh, Max Planck Institute for Informatics, Germany

3:15-3:35 New Lower Bounds for the Number of Pseudoline Arrangements

Adrian Dumitrescu and *Ritankar Mandal*, University of Wisconsin, Milwaukee, U.S.

3:40-4:00 Extremal and Probabilistic Results for Order Types

Jie Han, University of Rhode Island, U.S.; Yoshiharu Kohayakawa, Universidade de Sao Paulo, Brazil; *Marcelo Sales,* Emory University, U.S.; Henrique Stagni, Universidade of São Paulo, Brazil

Sunday, January 6

CP8

SODA Session 2C

2:00 PM-4:05 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

2:00-2:20 An Algorithmic Blend of LPs and Ring Equations for Promise CSPs

Joshua Brakensiek, Stanford University, U.S.; Venkatesan Guruswami, Carnegie Mellon University, U.S.

2:25-2:45 The Complexity of the Ideal Membership Problem for Constrained Problems Over the Boolean Domain

Monaldo Mastrolilli, IDSIA, Switzerland

2:50-3:10 Perfect Matchings, Rank of Connection Tensors and Graph Homomorphisms

Jin-Yi Cai and *Artsiom Hovarau*, University of Wisconsin, Madison, U.S.

3:15-3:35 Probabilistic Tensors and Opportunistic Boolean Matrix Multiplication

Petteri Kaski and Matti Karppa, Aalto University, Finland

3:40-4:00 Fast Greedy for Linear Matroids

Huy L. Nguyen, Northeastern University, U.S.

Coffee Break

4:05 PM-4:30 PM

Room:Ballroom Foyer - 2nd Floor

Sunday, January 6

CP9

ANALCO Session 3

4:30 PM-6:35 PM

Room: Diamond 2 - 2nd Floor

Chair: Markus Nebel, Universität Bielefeld, Germany

4:30-4:50 When Does Hillclimbing Fail on Monotone Functions?

Anders Martinsson, ETH Zürich, Switzerland

4:55-5:15 Degree Distributions of Generalized Hooking Networks

Colin Desmarais and Cecilia Holmgren, Uppsala University, Sweden

5:20-5:40 Subcritical Random Hypergraphs, High-Order Components, and Hypertrees

Wenjie Fang, Oliver Cooley, *Nicola Del Giudice*, and Mihyun Kang, Technische Universität, Graz, Austria

5:45-6:05 Random Walks on Graphs: New Bounds on Hitting, Meeting, Coalescing and Returning

Roberto Oliveira, IMPA, Rio de Janeiro, Brazil; Yuval Peres, Microsoft Research, U.S.

6:10-6:30 Arithmetic Progression Hypergraphs: Examining the Second Moment Method

Michael Mitzenmacher, Harvard University, U.S.

Sunday, January 6

CP10

SODA Session 3A

4:30 PM-6:35 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

4:30-4:50 Flow-Cut Gaps and Face Covers in Planar Graphs

Havana Rika and Robert Krauthgamer, Weizmann Institute of Science, Israel; James R. Lee, University of Washington, U.S.

4:55-5:15 On Constant Multi-Commodity Flow-Cut Gaps for Families of Directed Minor-Free Graphs

Ario Salmasi, Ohio State University, U.S.; Anastasios Sidiropoulos, University of Illinois at Chicago, U.S.; Vijay Sridhar, MathWorks, U.S.

5:20-5:40 Maximum Integer Flows in Directed Planar Graphs with Vertex Capacities and Multiple Sources and Sinks

Yipu Wang, University of Illinois at Urbana-Champaign, U.S.

5:45-6:05 A Faster Algorithm for Minimum-Cost Bipartite Matching in Minor-Free Graphs

Nathaniel Lahn and Sharath Raghvendra, Virginia Tech, U.S.

6:10-6:30 Finding Maximal Sets of Laminar 3-Separators in Planar Graphs in Linear Time

David Eppstein, University of California, Irvine, U.S.

Sunday, January 6

CP11

SODA Session 3B

4:30 PM-6:35 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

4:30-4:50 Reproducibility and Pseudo-determinism in Log-space

Yang Liu, Stanford University, U.S.; *Ofer Grossman*, University of California, Berkeley, U.S.

4:55-5:15 Pseudorandomness for Read-*k* DNF Formulas

Rocco A. Servedio, Columbia University, U.S.; *Li-Yang Tan*, Stanford University, U.S.

5:20-5:40 Near-optimal Bootstrapping of Hitting Sets for Algebraic Circuits

Mrinal Kumar, Simons Institute for the Theory of Computing, U.S.; Ramprasad Saptharishi and *Anamay Tengse*, Tata Institute of Fundamental Research, India

5:45-6:05 A Deterministic PTAS for the Algebraic Rank of Bounded Degree Polynomials

Vishwas Bhargava, Rutgers University, U.S.; Markus Bläser, Saarland University, Germany; *Gorav Jindal*, Aalto University, Finland; Anurag Pandey, Max Planck Institute for Informatics, Germany

6:10-6:30 Quantum Algorithms and Approximating Polynomials for Composed Functions with Shared Inputs

Mark Bun, Simons Institute for the Theory of Computing, U.S.; Robin Kothari, Microsoft Research, U.S.; Justin Thaler, Georgetown University, U.S.

Sunday, January 6

CP12

SODA Session 3C

4:30 PM-6:35 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

4:30-4:50 Anaconda: A Non-adaptive Conditional Sampling Algorithm for Distribution Testing

Gautam Kamath, Simons Institute for the Theory of Computing, U.S.; Christos Tzamos, University of Wisconsin, Madison, U.S.

4:55-5:15 Testing Halfspaces over Rotation-invariant Distributions

Nathaniel Harms, University of Waterloo, Canada

5:20-5:40 Every Testable (Infinite) Property of Bounded-degree Graphs Contains An Infinite Hyperfinite Subproperty

Hendrik Fichtenberger, Technische Universität Dortmund, Germany; Pan Peng, University of Sheffield, United Kingdom; Christian Sohler, Technische Universität Dortmund, Germany

5:45-6:05 Testing Matrix Rank, Optimally

Maria-Florina Balcan, Carnegie Mellon University, U.S.; Yi Li, Nanyang Technological University, Singapore; David Woodruff and *Hongyang Zhang*, Carnegie Mellon University, U.S.

6:10-6:30 A Ptas for ℓp-Low Rank Approximation

Frank Ban, University of California, Berkeley, U.S.; Vijay Bhattiprolu, Carnegie Mellon University, U.S.; Karl Bringmann and Pavel Kolev, Max Planck Institute for Informatics, Germany; Euiwoong Lee, New York University, U.S.; David Woodruff, Carnegie Mellon University, U.S.

Intermission

6:35 PM-6:45 PM

ALENEX/ANALCO Business Meeting

6:45 PM-7:45 PM

Room: Diamond 2 - 2nd Floor

Monday, January 7

Registration

8:00 AM-5:00 PM

Room: Registration Desk - 2rd Floor

Continental Breakfast

8:30 AM-9:00 AM

Room:Ballroom Foyer - 2nd Floor

Monday, January 7

CP13

ALENEX Session 1

9:00 AM-10:40 AM

Room: Diamond 2 - 2nd Floor

Chair: Henning Meyerhenke, Humboldt-Universität zu Berlin, Germany

9:00-9:20 Worst-case Efficient Sorting with QuickMergesort

Armin Weiß, Universität Stuttgart, Germany; Stefan Edelkamp, King's College London, United Kingdom

9:25-9:45 Simple and Fast BlockQuicksort using Lomuto's Partitioning Scheme

Martin Aumüller and *Nikolaj Hass*, IT University of Copenhagen, Denmark

9:50-10:10 Lightweight Distributed Suffix Array Construction

Florian Kurpicz and Johannes Fischer, Technische Universität Dortmund, Germany

10:15-10:35 Approximation of Trees by Self-nested Trees

Romain Azais, Inria, France; Jean-Baptiste Durand, Université Grenoble Alpes, France; Christophe Godin, Inria, France

Monday, January 7

CP14

SODA Session 4A

9:00 AM-11:05 AM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

9:00-9:20 Sublinear Algorithms for (Δ + 1) Vertex Coloring

Sepehr Assadi, Yu Chen, and Sanjeev Khanna, University of Pennsylvania, U.S.

9:25-9:45 Optimal Distributed Coloring Algorithms for Planar Graphs in the Local Model

Shiri Chechik and Doron Mukhtar, Tel Aviv University, Israel

9:50-10:10 Distributed Maximal Independent Set Using Small Messagesu

Mohsen Ghaffari, ETH Zürich, Switzerland

10:15-10:35 Distributed Triangle Detection via Expander Decomposition

Yi-Jun Chang and Seth Pettie, University of Michigan, U.S.; *Hengjie Zhang*, Tsinghua University, China

10:40-11:00 Oblivious Resampling Oracles and Parallel Algorithms for the Lopsided Lovász Local Lemma

David G. Harris, University of Maryland, U.S.

Monday, January 7

CP15

SODA Session 4B

9:00 AM-11:05 AM

Room: Topaz - 2nd Floor

Chair: To Be Determined

9:00-9:20 On the Number of Circuits in Regular Matroids (with Connections to Lattices and Codes)

Rohit Gurjar, Indian Institute of Technology Bombay, India; Nisheeth K. Vishnoi, École Polytechnique Fédérale de Lausanne, Switzerland

9:25-9:45 Minimum Cut and Minimum *k*-Cut in Hypergraphs via Branching Contractions

Kyle Fox, University of Texas, Dallas, U.S.; Debmalya Panigrahi, Duke University, U.S.; *Fred Zhang*, Harvard University, U.S.

9:50-10:10 Improving the Smoothed Complexity of Flip for Max-cut Problems

Ali Bibak, University of Illinois at Chicago, U.S.; Charles A. Carlson, University of Colorado Boulder, U.S.; Karthekeyan Chandrasekaran, University of Illinois at Urbana-Champaign, U.S.

10:15-10:35 Computing all Wardrop Equilibria Parametrized by the Flow Demand

Philipp Warode and Max Klimm, Humboldt University at Berlin, Germany

10:40-11:00 Nash Flows over Time with Spillback

Leon Sering, Technische Universität Berlin, Germany; Laura Vargas Koch, RWTH Aachen University, Germany

Monday, January 7

CP16

SODA Session 4C

9:00 AM-11:05 AM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

9:00-9:20 On the Rank of a Random Sparse Binary Matrix

Alan Frieze, Carnegie Mellon University, U.S.; Colin Cooper, King's College London, United Kingdom; Wesley Pegden, Carnegie Mellon University, U.S.

9:25-9:45 On Coalescence Time in Graphs-When Is Coalescing as Fast as Meeting?

Varun Kanade, University of Oxford, United Kingdom; *Frederik Mallmann-Trenn*, Massachusetts Institute of Technology, U.S.; Thomas Sauerwald, University of Cambridge, United Kingdom

9:50-10:10 Rapid Mixing of the Switch Markov Chain for Strongly Stable Degree Sequences and 2-Class Joint Degree Matrices

Georgios Amanatidis and *Pieter Kleer*, Centrum voor Wiskunde en Informatica (CWI), Netherlands

10:15-10:35 Xor Codes and Learning Sparse Parities with Noise

Andrej Bogdanov, Chinese University of Hong Kong, Hong Kong; *Manuel Sabin* and Prashant Nalini Vasudevan, University of California, Berkeley, U.S.

10:40-11:00 Seeded Graph Matching via Large Neighborhood Statistics

Elchanan Mossel, Massachusetts Institute of Technology, U.S.; *Jiaming Xu*, Duke University, U.S.

Coffee Break

11:05 AM-11:30 AM

Room: Ballroom Foyer - 2nd Floor

Monday, January 7

IP2

Recent Advances on the Complexity of Parameterized Counting Problems

11:30 AM-12:30 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

Research in the past few decades revealed that a many of the basic algorithmic problems are fixed-parameter tractable with various parameterizations: they can be solved in time f(k)n^c, where k is some parameter of the input instance. There have been significant progress both in the design of parameterized algorithms and in understanding the limitations of these techniques. However, up until recent years, there were relatively few algorithmic and complexity results on parameterized counting problems. Similarly to what one experiences in the area of polynomial-time computations, there are natural parameterized problems where finding a single solution is fixed-parameter tractable, but counting the number of solutions is hard. In this talk, I will give a survey of recent developments, including new results that give a clean and unified explanation for the complexity of #k-Path, #k-Matching, and many other parameterized counting problems.

Dániel Marx

Hungarian Academy of Sciences, Hungary

Lunch Break

12:30 PM-2:00 PM

Attendees on their own

Monday, January 7

CP17

ALENEX Session 2

2:00 PM-3:40 PM

Room: Diamond 2 - 2nd Floor

Chair: Andrew Goldberg, Amazon.com, U.S.

2:00-2:20 Fast and Exact Public Transit Routing with Restricted Pareto Sets
Daniel Delling, Julian Dibbelt, and *Thomas Pajor*, Apple, Inc., U.S.

2:25-2:45 Alternative Multicriteria Routes
Florian Barth and Stefan Funke, Universität Stuttgart, Germany; Sabine Storandt, Universität Konstanz, Germany

2:50-3:10 Concatenated *k*-Path Covers
Moritz Beck, Universität Konstanz, Germany; Kam-Yiu Lam, City University of Hong Kong, Hong Kong; Joseph Kee Yin Ng, Hong Kong Baptist University, Hong Kong; Sabine Storandt, Universität Konstanz, Germany; Chun Jiang Zhu, University of Connecticut, U.S.

3:15-3:35 Batch-Parallel Euler Tour Trees
Thomas Tseng, Laxman Dhulipala, and Guy Blelloch, Carnegie Mellon University, U.S.

Monday, January 7

CP18

SODA Session 5A

2:00 PM-4:05 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

2:00-2:20 Nearly ETH-tight Algorithms for Planar Steiner Tree with Terminals on Few Faces
Sándor Kisfaludi-Bak and *Jesper Nederlof*, Eindhoven University of Technology, Netherlands; Erik Jan van Leeuwen, Utrecht University, The Netherlands

2:25-2:45 Contraction Decomposition in Unit Disk Graphs and Algorithmic Applications in Parameterized Complexity
Meirav Zehavi, Ben-Gurion University, Israel; Fahad Panolan, University of Bergen, Norway; Saket Saurabh, Institute of Mathematical Sciences, India and University of Bergen, Norway

2:50-3:10 Polynomial-time Approximation Scheme for Minimum *k*-Cut in Planar and Minor-free Graphs
Alireza Farhadi, University of Maryland, U.S.; MohammadHossein Bateni, Google Research, U.S.; MohammadTaghi Hajiaghayi, University of Maryland, College Park, U.S.

3:15-3:35 Embedding Planar Graphs into Low-treewidth Graphs with Applications to Efficient Approximation Schemes for Metric Problems
Eli Fox-Epstein, Tufts University, U.S.; *Philip Klein*, Brown University, U.S.; Aaron Schild, University of California, Berkeley, U.S.

3:40-4:00 A PTAS for Euclidean TSP with Hyperplane Neighborhoods
Antonios Antoniadis and Krzysztof Fleszar, Max Planck Institute for Informatics, Germany; Ruben Hoeksma, Universität Bremen, Germany; Kevin Schewior, Technische Universität München, Germany and École Normale Supérieure, France

Monday, January 7

CP19

SODA Session 5B

2:00 PM-4:05 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

2:00-2:20 The Streaming *k*-Mismatch Problem
Raphael Clifford, University of Bristol, United Kingdom; *Tomasz Kociumaka*, University of Warsaw, Poland; Ely Porat, Bar-Ilan University, Israel

2:25-2:45 Few Matches or Almost Periodicity: Faster Pattern Matching with Mismatches in Compressed Texts
Karl Bringmann, Marvin Künnemann, and *Philip Wellnitz*, Max Planck Institute for Informatics, Germany

2:50-3:10 Lower Bounds for Text Indexing with Mismatches and Differences
Tatiana Starikovskaya, École Normale Supérieure Paris, France; Vincent Cohen-Addad, Sorbonne Universités and CNRS, France; Laurent Feuilloley, Universite Paris Diderot, France

3:15-3:35 Efficiently Approximating Edit Distance Between Pseudorandom Strings
William Kuszmaul, Massachusetts Institute of Technology, U.S.

3:40-4:00 Approximating Lcs in Linear Time: Beating the $\sqrt{n}$ Barrier
Saeed Seddighin and MohammadTaghi Hajiaghayi, University of Maryland, College Park, U.S.; Masoud Seddighin, Sharif University of Technology, Iran; Xiaorui Sun, University of Illinois at Chicago, U.S.

Monday, January 7

CP20

SODA Session 5C

2:00 PM-4:05 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

2:00-2:20 The Maximum Number of Minimal Dominating Sets in a Tree
Günter Rote, Freie Universität Berlin, Germany

2:25-2:45 How to Guess an n-Digit Number
Zilin Jiang, Massachusetts Institute of Technology, U.S.; *Nikita Polyanskii*, Skolkovo Institute of Science and Technology, Russia

2:50-3:10 Vector Clique Decompositions
Raphael Yuster, University of Haifa, Israel

3:15-3:35 Four-Coloring P_6-Free Graphs
Maria Chudnovsky and Sophie Spirkl, Princeton University, U.S.; *Mingxian Zhong*, Lehman College, CUNY, U.S.

3:40-4:00 Polynomial-time Algorithm for Maximum Weight Independent Set on P_6-Free Graphs
Michal Pilipczuk, University of Warsaw, Poland; Andrzej Grzesik, Jagiellonian University, Poland; Tereza Klimošová, Charles University, Czech Republic; Marcin Pilipczuk, University of Warsaw, Poland

Coffee Break

4:05 PM-4:30 PM

Room:Ballroom Foyer - 2nd Floor

Monday, January 7

CP21

ALENEX Session 3

4:30 PM-6:10 PM

Room: Diamond 2 - 2nd Floor

Chair: Cliff Stein, Columbia University, U.S.

4:30-4:50 A New Integer Linear Program for the Steiner Tree Problem with Revenues, Budget and Hop Constraints
Adalat Jabrayilov and Petra Mutzel, Technische Universität Dortmund, Germany

4:55-5:15 SAT-Encodings for Treecut Width and Treedepth
Robert Ganian and *Neha Lodha*, Technische Universität Wien, Austria; Sebastian Ordyniak, University of Sheffield, United Kingdom; Stefan Szeider, Technische Universität Wien, Austria

5:20-5:40 Efficiently Enumerating Hitting Sets of Hypergraphs Arising in Data Profiling
Thomas Bläsius, Tobias Friedrich, and Julius Lischeid, Universität Potsdam, Germany; Kitty Meeks, University of Glasgow, Scotland, United Kingdom; *Martin Schirneck*, Universität Potsdam, Germany

5:45-6:05 Exactly Solving the Maximum Weight Independent Set Problem on Large Real-world Graphs
Darren Strash, Colgate University, U.S.; *Sebastian Lamm*, Karlsruhe Institute of Technology, Germany; Christian Schulz, University of Vienna, Austria and Karlsruhe Institute of Technology, Germany; Robert Williger, Karlsruhe Institute of Technology, Germany; Huashuo Zhang, Colgate University, U.S.

Monday, January 7

CP22

SODA Session 6A

4:30 PM-6:35 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

4:30-4:50 Strategies for Stable Merge Sorting
Samuel Buss and *Alexander Knop*, University of California, San Diego, U.S.

4:55-5:15 A Sort of an Adversary
Or Zamir, Haim Kaplan, and Uri Zwick, Tel Aviv University, Israel

5:20-5:40 A New Path from Splay to Dynamic Optimality
Caleb Levy, Princeton University, U.S.; Robert Tarjan, Princeton University and Intertrust Technologies, U.S.

5:45-6:05 A Faster External Memory Priority Queue with DecreaseKeys
Shunhua Jiang, Tsinghua University, China; Kasper G. Larsen, Aarhus University, Denmark

6:10-6:30 Optimal Construction of Compressed Indexes for Highly Repetitive Texts
Dominik Kempa, University of Warwick, United Kingdom

Monday, January 7

CP23

SODA Session 6B

4:30 PM-6:35 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

4:30-4:50 Approximability of P -> Q Matrix Norms: Generalized Krivine Rounding and Hypercontractive Hardness

Vijay Bhattiprolu, Carnegie Mellon University, U.S.; Mrinalkanti Ghosh, Toyota Technological Institute at Chicago, U.S.; Venkatesan Guruswami, Carnegie Mellon University, U.S.; Euiwoong Lee, New York University, U.S.; Madhur Tulsiani, Toyota Technological Institute at Chicago, U.S.

4:55-5:15 Proportional Volume Sampling and Approximation Algorithms for A-Optimal Design

Uthaipon Tantipongpipat and Mohit Singh, Georgia Institute of Technology, U.S.; Aleksandar Nikolov, University of Toronto, Canada

5:20-5:40 Perron-Frobenius Theory in Nearly Linear Time: Positive Eigenvectors, M-Matrices, Graph Kernels, and Other Applications

Amirmahdi Ahmadinejad, Arun Jambulapati, Amin Saberi, and Aaron Sidford, Stanford University, U.S.

5:45-6:05 Iterative Refinement for ℓp-Norm Regression

Sushant Sachdeva and *Deeksha Adil*, University of Toronto, Canada; Rasmus Kyng, Yale University, U.S.; Richard Peng, Georgia Institute of Technology, U.S.

6:10-6:30 Optimizing Quantum Optimization Algorithms via Faster Quantum Gradient Computation

Andras Gilyen, QuSoft, CWI and University of Amsterdam, Netherlands; Srinivasan Arunachalam, Massachusetts Institute of Technology, U.S.; Nathan Wiebe, Microsoft Research, U.S.

Monday, January 7

CP24

SODA Session 6C

4:30 PM-6:35 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

4:30-4:50 Towards Tight(er) Bounds for the Excluded Grid Theorem

Julia Chuzhoy, Toyota Technological Institute at Chicago, U.S.; *Zihan Tan*, University of Chicago, U.S.

4:55-5:15 Polynomial Planar Directed Grid Theorem

Meike Hatzel, Technische Universität Berlin, Germany; Ken-ichi Kawarabayashi, National Institute of Informatics, Japan; Stephan Kreutzer, Technische Universität Berlin, Germany

5:20-5:40 A Tight Erdös-Pósa Function for Planar Minors

Jean-Florent Raymond, Technische Universität Berlin, Germany; Wouter Cames Van Batenburg, *Tony Huynh*, and Gwenaël Joret, Université Libre de Bruxelles, Belgium

5:45-6:05 Polynomial Bounds for Centered Colorings on Proper Minor-Closed Graph Classes

Sebastian Siebertz, Humboldt University at Berlin, Germany; Michal Pilipczuk, University of Warsaw, Poland

6:10-6:30 Every Collinear Set in a Planar Graph is Free

Vida Dujmovic, University of Ottawa, Canada; Fabrizio Frati, Universita Roma Tre, Italy; Daniel Gonçalves, Université de Montpellier and CNRS, France; Pat Morin, Carleton University, Canada; *Günter Rote*, Freie Universität Berlin, Germany

Intermission

6:35 PM-6:45 PM

SODA Business Meeting and Awards Presentation

6:45 PM-7:45 PM

Room: Emerald Ballroom - 2nd Floor

Tuesday, January 8

Registration

8:00 AM-5:00 PM

Room: Registration Desk - 2rd Floor

Continental Breakfast

8:30 AM-9:00 AM

Room: Ballroom Foyer - 2nd Floor

Tuesday, January 8

CP25

ALENEX Session 4

9:00 AM-11:05 AM

Room: Diamond 2 - 2nd Floor

Chair: Stephen Kobourov, University of Arizona, U.S.

9:00-9:20 Parallel Range, Segment and Rectangle Queries with Augmented Maps
Yihan Sun and Guy Blelloch, Carnegie Mellon University, U.S.

9:25-9:45 A Practical Algorithm for Spatial Agglomerative Clustering
Thom Castermans, Bettina Speckmann, and Kevin Verbeek, Eindhoven University of Technology, Netherlands

9:50-10:10 Practical Methods for Computing Large Covering Tours and Cycle Covers with Turn Cost
Sándor Fekete and *Dominik M. Krupke*, Technische Universität Braunschweig, Germany

10:15-10:35 Faster Support Vector Machines
Sebastian Schlag and Matthias Schmitt, Karlsruhe Institute of Technology, Germany; Christian Schulz, University of Vienna, Austria and Karlsruhe Institute of Technology, Germany

10:40-11:00 Scalable Edge Partitioning
Darren Strash, Colgate University, U.S.; Sebastian Schlag, Karlsruhe Institute of Technology, Germany; Christian Schulz, University of Vienna, Austria and Karlsruhe Institute of Technology, Germany; *Daniel Seemaier*, Karlsruhe Institute of Technology, Germany

Tuesday, January 8

CP26

SODA Session 7A

9:00 AM-11:05 AM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

9:00-9:20 A 1.5-Approximation for Path TSP
Rico Zenklusen, ETH Zürich, Switzerland

9:25-9:45 A New Dynamic Programming Approach for Spanning Trees with Chain Constraints and Beyond
Martin Nagele and Rico Zenklusen, ETH Zürich, Switzerland

9:50-10:10 Lift and Project Algorithms for Precedence Constrained Scheduling to Minimize Completion Time
Janardhan Kulkarni, Microsoft Research, U.S.; Shashwat Garg, Eindhoven University of Technology, Netherlands; Shi Li, State University of New York at Buffalo, U.S.

10:15-10:35 A Polynomial Time Constant Approximation For Minimizing Total Weighted Flow-time
Janardhan Kulkarni, Microsoft Research, U.S.; Uriel Feige, Weizmann Institute of Science, Israel; Shi Li, State University of New York at Buffalo, U.S.

10:40-11:00 On Approximating (Sparse) Covering Integer Programs
Chandra Chekuri and *Kent Quanrud*, University of Illinois at Urbana-Champaign, U.S.

Tuesday, January 8

CP27

SODA Session 7B

9:00 AM-11:05 AM

Room: Topaz - 2nd Floor

Chair: To Be Determined

9:00-9:20 Coresets Meet EDCS: Algorithms for Matching and Vertex Cover on Massive Graphs
Sepehr Assadi, University of Pennsylvania, U.S.; MohammadHossein Bateni, Google Research, U.S.; Aaron Bernstein, Rutgers University, U.S.; Vahab Mirrokni, Google, Inc., U.S.; Cliff Stein, Columbia University, U.S.

9:25-9:45 Sparsifying Distributed Algorithms with Ramifications in Massively Parallel Computation and Centralized Local Computation
Mohsen Ghaffari and Jara Uitto, ETH Zürich, Switzerland

9:50-10:10 Massively Parallel Approximation Algorithms for Edit Distance and Longest Common Subsequence
Saeed Seddighin, University of Maryland, College Park, U.S.; Xiaorui Sun, University of Illinois at Chicago, U.S.; MohammadTaghi Hajiaghayi, University of Maryland, College Park, U.S.

10:15-10:35 Low Congestion Cycle Covers and Their Applications
Eylon Yogev and Merav Parter, Weizmann Institute of Science, Israel

10:40-11:00 Distributed Algorithms Made Secure: A Graph Theoretic Approach
Eylon Yogev and Merav Parter, Weizmann Institute of Science, Israel

Tuesday, January 8

CP28

SODA Session 7C

9:00 AM-11:05 AM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

9:00-9:20 Interval Vertex Deletion Admits a Polynomial Kernel

Akanksha Agrawal, Hungarian Academy of Sciences, Hungary

9:25-9:45 Losing Treewidth by Separating Subsets

Euiwoong Lee, New York University, U.S.; Anupam Gupta and Jason M. Li, Carnegie Mellon University, U.S.; Pasin Manurangsi, University of California, Berkeley, U.S.; *Michal Wlodarczyk*, University of Warsaw, Poland

9:50-10:10 On *r*-Simple *k*-Path and Related Problems Parameterized by *k*/*r*

Meirav Zehavi, Ben-Gurion University, Israel; Gregory Gutin and Magnus Wahlstrom, Royal Holloway, University of London, United Kingdom

10:15-10:35 A Time and Space-Optimal Algorithm for the Many-Visits TSP

André Berger, Maastricht University, Netherlands; László Kozma, Eindhoven University of Technology, Netherlands; Matthias Mnich, Maastricht University, Netherlands and Universität Bonn, Germany; *Roland Vincze*, Maastricht University, Netherlands

10:40-11:00 Quantum Speedups for Exponential-Time Dynamic Programming Algorithms

Jevgenijs Vihrovs, Andris Ambainis, Kaspars Balodis, Janis Iraids, Martins Kokainis, and *Krisjanis Prusis*, University of Latvia, Latvia

Coffee Break

11:05 AM-11:30 AM

Room: Ballroom Foyer - 2nd Floor

Tuesday, January 8

IP3

Inherent Trade-Offs in Algorithmic Fairness

11:30 AM-12:30 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

Recent discussion in the public sphere about classification by algorithms has involved tension between competing notions of what it means for such a classification to be fair to different groups. We consider several of the key fairness conditions that lie at the heart of these debates, and discuss recent research establishing inherent trade-offs between these conditions. We also consider a variety of methods for promoting fairness and related notions for classification and selection problems that involve sets rather than just individuals. This talk is based on joint work with Sendhil Mullainathan, Manish Raghavan, and Maithra Raghu.

Jon M. Kleinberg

Cornell University, U.S.

Lunch Break

12:30 PM-2:00 PM

Attendees on their own

Tuesday, January 8

CP29

SOSA Session 1

2:00 PM-4:05 PM

Room: Diamond 2 - 2nd Floor

Chair: To Be Determined

2:00-2:20 Isotonic Regression by Dynamic Programming

Günter Rote, Freie Universität Berlin, Germany

2:25-2:45 An Illuminating Algorithm for the Light Bulb Problem

Josh Alman, Massachusetts Institute of Technology, U.S.

2:50-3:10 Simple Concurrent Labeling Algorithms for Connected Components

Sixue Liu, Princeton University, U.S.; Robert Tarjan, Princeton University and Intertrust Technologies, U.S.

3:15-3:35 A Framework for Searching in Graphs in the Presence of Errors

Dariusz Dereniowski, Gdansk University of Technology, Poland; Stefan Tiegel, Przemyslaw Uznanski, and Daniel Wolleb-Graf, ETH Zürich, Switzerland

3:40-4:00 Selection from Heaps, Row-Sorted Matrices and X+Y Using Soft Heaps

Uri Zwick, Haim Kaplan, and *Or Zamir*, Tel Aviv University, Israel

Tuesday, January 8

CP30

SODA Session 8A

2:00 PM-4:05 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

2:00-2:20 Optimal Las Vegas Approximate Near Neighbors in ℓp
Alexander Wei, Harvard University, U.S.

2:25-2:45 The Andoni-Krauthgamer-Razenshteyn Characterization of Sketchable Norms Fails for Sketchable Metrics
Assaf Naor, Princeton University, U.S.; Subhash Khot, New York University, U.S.

2:50-3:10 Tight Bounds for ℓp Oblivious Subspace Embeddings
Ruosong Wang and David Woodruff, Carnegie Mellon University, U.S.

3:15-3:35 Optimal Lower Bounds for Distributed and Streaming Spanning Forest Computation
Huacheng Yu and Jelani Nelson, Harvard University, U.S.

3:40-4:00 Communication-Rounds Tradeoffs for Common Randomness and Secret Key Generation
Mitali Bafna, Harvard University, U.S.; Badih Ghazi, Google Research, U.S.; *Noah Golowich* and Madhu Sudan, Harvard University, U.S.

Tuesday, January 8

CP31

SODA Session 8B

2:00 PM-4:05 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

3:40-4:00 Deterministically Maintaining a (2 +ϵ)-Approximate Minimum Vertex Cover in $O(1/\varepsilon^2)$ Amortized Update Time
Sayan Bhattacharya, University of Warwick, United Kingdom; Janardhan Kulkarni, Microsoft Research, U.S.

2:00-2:20 (1 + Eps)-Approximate Incremental Matching in Constant Deterministic Amortized Time
Chris Schwiegelshohn, Università di Roma "La Sapienza", Italy; Fabrizio Grandoni, IDSIA, Switzerland; Stefano Leonardi, Università di Roma "La Sapienza", Italy; Piotr Sankowski, University of Warsaw, Poland; Shay Solomon, Tel Aviv University, Israel

2:25-2:45 A Deamortization Approach for Dynamic Spanner and Dynamic Maximal Matching
Aaron Bernstein, Rutgers University, U.S.; Sebastian Forster, University of Salzburg, Germany; Henzinger Monika, University of Vienna, Austria

2:50-3:10 Fully Dynamic Maximal Independent Set with Sublinear in n Update Time
Sepehr Assadi, University of Pennsylvania, U.S.; Krzysztof Onak, IBM Research, U.S.; Baruch Schieber, New Jersey Institute of Technology, U.S.; *Shay Solomon*, Tel Aviv University, Israel

3:15-3:35 Dynamic Edge Coloring with Improved Approximation
Ran Duan, Haoqing He, and *Tianyi Zhang*, Tsinghua University, China

Tuesday, January 8

CP32

SODA Session 8C

2:00 PM-4:05 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

2:00-2:20 Prophet Secretary Through Blind Strategies
José Correa and *Raimundo Saona*, Universidad de Chile, Chile; Bruno Ziliotto, Universite Paris Dauphine and CNRS, France

2:25-2:45 Pricing for Online Resource Allocation: Intervals and Paths
Shuchi Chawla, Benjamin Miller, and *Yifeng Teng*, University of Wisconsin, Madison, U.S.

2:50-3:10 An Optimal Truthful Mechanism for the Online Weighted Bipartite Matching Problem
Rebecca Reiffenhaeuser, Università di Roma "La Sapienza", Italy

3:15-3:35 Fully Polynomial-Time Approximation Schemes for Fair Rent Division
Eshwar Ram Arunachaleswaran, Siddharth Barman, and *Nidhi Rathi*, Indian Institute of Science, India

3:40-4:00 Communication Complexity of Discrete Fair Division
Benjamin Plaut and Tim Roughgarden, Stanford University, U.S.

Coffee Break

4:05 PM-4:30 PM

Room: Ballroom Foyer - 2nd Floor

Tuesday, January 8

CP33

SOSA Session 2

4:30 PM-6:35 PM

Room: Diamond 2 - 2nd Floor

Chair: To Be Determined

4:30-4:50 Approximating Optimal Transport with Linear Programs
Kent Quanrud, University of Illinois at Urbana-Champaign, U.S.

4:55-5:15 LP Relaxation and Tree Packing for Minimum *k*-Cuts
Chandra Chekuri and Kent Quanrud, University of Illinois at Urbana-Champaign, U.S.; *Chao Xu*, Yahoo! Research, U.S.

5:20-5:40 On Primal-Dual Circle Representations
Stefan Felsner, Technische Universität, Berlin, Germany; *Günter Rote*, Freie Universität Berlin, Germany

5:45-6:05 Asymmetric Convex Intersection Testing
Luis Barba, ETH Zürich, Switzerland; Wolfgang J. Mulzer, Freie Universität Berlin, Germany

6:10-6:30 Relaxed Voronoi: A Simple Framework for Terminal-Clustering Problems
Arnold Filtser, Ben Gurion University, Israel; Robert Krauthgamer and Ohad Trabelsi, Weizmann Institute of Science, Israel

Tuesday, January 8

CP34

SODA Session 9A

4:30 PM-6:35 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

4:30-4:50 The I/O Complexity of Toom-Cook Integer Multiplication
Lorenzo De Stefani, Brown University, U.S.; Gianfranco Bilardi, University of Padova, Italy

4:55-5:15 I/O-Efficient Algorithms for Topological Sort and Related Problems
Nairen Cao, Jeremy Fineman, *Katina Russell*, and Eugene Yang, Georgetown University, U.S.

5:20-5:40 On the Structure of Unique Shortest Paths in Graphs
Greg Bodwin, Georgia Institute of Technology, U.S.

5:45-6:05 Near Optimal Algorithms For The Single Source Replacement Paths Problem
Sarel Cohen and *Shiri Chechik*, Tel Aviv University, Israel

6:10-6:30 Exact Distance Oracles for Planar Graphs with Failing Vertices
Panagiotis Charalampopoulos, King's College London, United Kingdom; Shay Mozes and Benjamin Tebeka, Interdisciplinary Center Herzliya, Israel

Tuesday, January 8

CP35

SODA Session 9B

4:30 PM-6:35 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

4:30-4:50 Analyzing Boolean Functions on the Biased Hypercube Via Higher-Dimensional Agreement Tests
Yuval Filmus, Technion Israel Institute of Technology, Israel; *Irit Dinur*, Weizmann Institute of Science, Israel; Prahladh Harsha, Tata Institute of Fundamental Research, India

4:55-5:15 List Decoding with Double Samplers
Inbal R. Livni Navon and Irit Dinur, Weizmann Institute of Science, Israel; Prahladh Harsha, Tata Institute of Fundamental Research, India; Tali Kaufman, Massachusetts Institute of Technology, U.S.; Amnon Ta-Shma, Tel Aviv University, Israel

5:20-5:40 Maximally Recoverable LRCs: A Field Size Lower Bound and Constructions for Few Heavy Parities
Sivakanth Gopi, Microsoft Research, U.S.; Venkatesan Guruswami, Carnegie Mellon University, U.S.; Sergey Yekhanin, Microsoft, U.S.

5:45-6:05 Binary Robust Positioning Patterns with Low Redundancy and Efficient Locating Algorithms
Yeow Meng Chee, Duc Tu Dao, Han Mao Kiah, San Ling, and *Hengjia Wei*, Nanyang Technological University, Singapore

6:10-6:30 Synchronization Strings: Highly Efficient Deterministic Constructions over Small Alphabets
Kuan Cheng, Johns Hopkins University, U.S.; Bernhard Haeupler, Carnegie Mellon University, U.S.; Xin Li, Johns Hopkins University, U.S.; Amirbehshad Shahrasbi, Carnegie Mellon University, U.S.; *Ke Wu*, Johns Hopkins University, U.S.

Tuesday, January 8

CP36

SODA Session 9C

4:30 PM-6:35 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

4:30-4:50 The Complexity of Approximately Counting Retractions
Jacob Focke, Leslie Ann Goldberg, and Stanislav Zivny, University of Oxford, United Kingdom

4:55-5:15 Improved Bounds for Randomly Sampling Colorings Via Linear Programming
Sitan Chen, Massachusetts Institute of Technology, U.S.; Michelle Delcourt, University of Waterloo, Canada; Ankur Moitra, Massachusetts Institute of Technology, U.S.; Guillem Perarnau, University of Birmingham, United Kingdom; Luke Postle, University of Waterloo, Canada

5:20-5:40 Algorithms for #{BIS}-Hard Problems on Expander Graphs
Matthew Jenssen and Peter Keevash, University of Oxford, United Kingdom; Will Perkins, University of Illinois at Chicago, U.S.

5:45-6:05 Approximability of the Six-Vertex Model
Jin-Yi Cai and *Tianyu Liu*, University of Wisconsin, Madison, U.S.; Pinyan Lu, Shanghai University of Finance and Economics, China

6:10-6:30 Zeros of Holant Problems: Locations and Algorithms
Heng Guo, University of Edinburgh, United Kingdom; Chao Liao, Shanghai Jiao Tong University, China; Pinyan Lu, Shanghai University of Finance and Economics, China; *Chihao Zhang*, Shanghai Jiaotong University, China

Intermission

6:35 PM-6:45 PM

SOSA Business Meeting

6:45 PM-7:45 PM

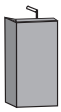

Room: Diamond 2 - 2nd Floor

Wednesday, January 9

Registration

8:00 AM-5:00 PM

Room: Registration Desk - 2rd Floor

Continental Breakfast

8:30 AM-9:00 AM

Room: Ballroom Foyer - 2nd Floor

Wednesday, January 9

CP37

SOSA Session 3

9:00 AM-11:05 AM

Room: Diamond 2 - 2nd Floor

Chair: To Be Determined

9:00-9:20 Towards a Unified Theory of Sparsification for Matching Problems
Sepehr Assadi, University of Pennsylvania, U.S.; Aaron Bernstein, Rutgers University, U.S.

9:25-9:45 A New Application of Orthogonal Range Searching for Computing Giant Graph Diameters
Guillaume Ducoffe, ICI – National Institute for Research and Development informatics, Romania

9:50-10:10 Simplified and Space-Optimal Semi-Streaming (2+ε)-Approximate Matching
Mohsen Ghaffari, ETH Zürich, Switzerland; *David Wajc*, Carnegie Mellon University, U.S.

10:15-10:35 Simple Greedy 2-Approximation Algorithm for the Maximum Genus of a Graph
Michal Kotrbcik and *Martin Skoviera*, Comenius University, Bratislava, Slovakia

10:40-11:00 A Note on Max K-Vertex Cover: Faster Fpt-As, Smaller Approximate Kernel and Improved Approximation
Pasin Manurangsi, University of California, Berkeley, U.S.

Wednesday, January 9

CP38

SODA Session 10A

9:00 AM-11:05 AM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

9:00-9:20 On Facility Location with General Lower Bounds

Shi Li, State University of New York at Buffalo, U.S.

9:25-9:45 Hierarchical Clustering Better Than Average-Linkage

Moses Charikar, *Vaggos Chatziafratis*, and Rad Niazadeh, Stanford University, U.S.

9:50-10:10 The Threshold for SDP-Refutation of Random Regular NAE-3SAT

Tselil Schramm, Massachusetts Institute of Technology and Harvard University, U.S.; Yash Deshpande, Massachusetts Institute of Technology, U.S.; Andrea Montanari, Stanford University, U.S.; Ryan O'Donnell, Carnegie Mellon University, U.S.; Subhabrata Sen, Massachusetts Institute of Technology, U.S.

10:15-10:35 Exponential Lower Bounds on Spectrahedral Representations of Hyperbolicity Cones

Nick Ryder, Nikhil Srivastava, and Prasad Raghavendra, University of California, Berkeley, U.S.; Benjamin Weitz, Pintrest, U.S.

10:40-11:00 Universal Trees Grow Inside Separating Automata: Quasi-Polynomial Lower Bounds for Parity Games

Wojciech Czerwinski, University of Warsaw, Poland; Laure Daviaud, University of Warwick, United Kingdom; Nathanael Fijalkow, Université Bordeaux, France; Marcin Jurdzinski and *Ranko Lazic*, University of Warwick, United Kingdom; Pawel Parys, University of Warsaw, Poland

Wednesday, January 9

CP39

SODA Session 10B

9:00 AM-11:05 AM

Room: Topaz - 2nd Floor

Chair: To Be Determined

9:00-9:20 Theorems of Carathéodory, Helly, and Tverberg Without Dimension

Karim Adiprasito, Hebrew University of Jerusalem, Israel; Imre Bárány, Alfréd Rényi Institute of Mathematics, Budapest and University College London, United Kingdom; *Nabil H. Mustafa*, ESIEE, France

9:25-9:45 On the Spanning and Routing Ratio of Theta-Four

Darryl R. Hill and Prosenjit Bose, Carleton University, Canada; Jean-Lou De Carufel, University of Ottawa, Canada; Michiel Smid, Carleton University, Canada

9:50-10:10 Greedy Spanners Are Optimal in Doubling Metrics

Glencora Borradaile and *Hung Le*, Oregon State University, U.S.; Christian Wulff-Nilsen, University of Copenhagen, Denmark

10:15-10:35 Viewing the Rings of a Tree: Minimum Distortion Embeddings into Trees

Benjamin A. Raichel, University of Texas at Dallas, U.S.; Amir Nayyeri, Oregon State University, U.S.

10:40-11:00 Minimizing Interference Potential Among Moving Entities

Daniel Busto, Ericsson, Montreal, Canada; *William S. Evans*, University of British Columbia, Canada; David Kirkpatrick, University of British Columbia, Canada

Wednesday, January 9

CP40

SODA Session 10C

9:00 AM-11:05 AM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

9:00-9:20 Can We Overcome the *n* log *n* Barrier for Oblivious Sorting?

Wei-Kai Lin and Elaine Shi, Cornell University, U.S.; *Tiancheng Xie*, Shanghai Jiao Tong University, China

9:25-9:45 Lower Bounds for Oblivious Data Structures

Riko Jacob, IT University of Copenhagen, Denmark; *Kasper G. Larsen* and Jesper B. Nielsen, Aarhus University, Denmark

9:50-10:10 Foundations of Differentially Oblivious Algorithms

T-H. Hubert Chan, University of Hong Kong, Hong Kong; Kai-Min Chung, Academia Sinica, Taiwan; Bruce M. Maggs, Duke University and Akamai Technologies, U.S.; *Elaine Shi*, Cornell University, U.S.

10:15-10:35 Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity

Ulfar Erlingsson, Vitaly Feldman, Ilya Mironov, *Ananth Raghunathan*, and Kunal Talwar, Google, Inc., U.S.; Abhradeep Thakurta, Google and University of California, Santa Cruz, U.S.

10:40-11:00 Towards Instance-Optimal Private Query Release

Aleksandar Nikolov, University of Toronto, Canada; *Jaroslaw Blasiok*, Harvard University, U.S.; Mark Bun, Princeton University, U.S.; Thomas Steinke, IBM Research, U.S.

Coffee Break

11:05 AM-11:30 AM

Room: Ballroom Foyer - 2nd Floor

Wednesday, January 9

IP4

(Hardness of) Approximation and Expansion

11:30 AM-12:30 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

Recently, there has been some exciting progress towards proving the unique games conjecture. Is the conjecture true? Experts are not sure. Still, it is clearly key to understanding approximation algorithms for many optimization problems, most notably constraint satisfaction. I will describe the background and consequences of the conjecture and then focus on some interesting technical step pertaining to the *expansion* of the so-called Grassman graph. I will then describe a more general phenomenon called high dimensional expansion (HDX). HDX is family of a generalizations of expansion in graphs to hypergraphs. This is a new area of study and has potential for algorithms, as I will try to show.

Irit Dinur
Weizmann Institute of Science, Israel

Lunch Break

12:30 PM-2:00 PM

Attendees on their own

Wednesday, January 9

CP41

SOSA Session 4

2:00 PM-4:05 PM

Room: Diamond 2 - 2nd Floor

Chair: To Be Determined

2:00-2:20 Simple Contention Resolution Via Multiplicative Weight Updates
Yi-Jun Chang, Wenyu Jin, and *Seth Pettie*, University of Michigan, U.S.

2:25-2:45 A Simple Near-Linear Pseudopolynomial Time Randomized Algorithm for Subset Sum
Ce Jin and Hongxun Wu, Tsinghua University, China

2:50-3:10 Submodular Optimization in the MapReduce Model
Paul Liu and Jan Vondrak, Stanford University, U.S.

3:15-3:35 Compressed Sensing with Adversarial Sparse Noise Via L1 Regression
Sushrut Karmalkar and Eric Price, University of Texas at Austin, U.S.

3:40-4:00 Approximating Maximin Share Allocations
Jugal Garg, Peter McGlaughlin, and *Setareh Taki*, University of Illinois at Urbana-Champaign, U.S.

Wednesday, January 9

CP42

SODA Session 11A

2:00 PM-4:05 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

2:00-2:20 Non-Empty Bins with Simple Tabulation Hashing
Anders Aamand and Mikkel Thorup, University of Copenhagen, Denmark

2:25-2:45 Derandomized Balanced Allocation
Xue Chen, Northwestern University, U.S.

2:50-3:10 Optimal Ball Recycling
Michael A. Bender and Jake Christensen, Stony Brook University, U.S.; *Alexander Conway* and Martin Farach-Colton, Rutgers University, U.S.; Rob Johnson, Vmware Research, U.S.; Meng-Tsung Tsai, National Chiao Tung University, Taiwan

3:15-3:35 A Fourier-Analytic Approach for the Discrepancy of Random Set Systems
Rebecca Hoberg and Thomas Rothvoss, University of Washington, U.S.

3:40-4:00 On the Discrepancy of Random Low Degree Set Systems
Nikhil Bansal, Centrum voor Wiskunde en Informatica (CWI) and Eindhoven University of Technology, Netherlands; *Raghu Meka*, University of California, Los Angeles, U.S.

Wednesday, January 9

CP43

SODA Session 11B

2:00 PM-4:05 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

2:00-2:20 Optimal Lower Bounds for Sketching Graph Cuts

Charles A. Carlson and Alexandra Kolla, University of Colorado Boulder, U.S.; Nikhil Srivastava and Luca Trevisan, University of California, Berkeley, U.S.

2:25-2:45 Spectral Sparsification of Hypergraphs

Tasuku Soma, University of Tokyo, Japan; Yuichi Yoshida, National Institute of Informatics, Japan

2:50-3:10 Cheeger Inequalities for Submodular Transformations

Yuichi Yoshida, National Institute of Informatics, Japan

3:15-3:35 Short Cycles via Low-diameter Decompositions

Yang Liu, Stanford University, U.S.; Sushant Sachdeva and Zejun Yu, University of Toronto, Canada

3:40-4:00 Expander Decomposition and Pruning: Faster, Stronger, and Simpler.

Thatchaphol Saranurak, Toyota Technological Institute at Chicago, U.S.; *Di Wang*, Georgia Institute of Technology, U.S.

Wednesday, January 9

CP44

SODA Session 11C

2:00 PM-4:05 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

2:00-2:20 Constructive Polynomial Partitioning for Algebraic Curves In~ \reals^3 with Applications

Esther Ezra, Georgia Institute of Technology, U.S.; Boris Aronov, New York University, U.S.; Joshua Zahl, University of British Columbia, Canada

2:25-2:45 Computing Height Persistence and Homology Generators in R^3 Efficiently

Tamal K. Dey, Ohio State University, U.S.

2:50-3:10 Hardness of Approximation for Morse Matching

Abhishek J. Rathod and Ulrich Bauer, Technische Universität München, Germany

3:15-3:35 Improved Topological Approximations by Digitization

Aruni Choudhary, Freie Universität Berlin, Germany; Michael Kerber, Graz University of Technology, Austria; Sharath Raghvendra, Virginia Tech, U.S.

3:40-4:00 Full Tilt: Universal Constructors for General Shapes with Uniform External Forces

Jose Balanza-Martinez, David Caballero, Angel Cantu, Luis Garcia, *Austin Luchsinger*, and Rene Reyes, University of Texas, Rio Grande Valley; Robert T. Schweller, University of Texas - Pan American, U.S.; Tim Wylie, University of Texas, Rio Grande Valley

Coffee Break

4:05 PM-4:30 PM

Room: Ballroom Foyer - 2nd Floor

Wednesday, January 9

CP45

SODA Session 12A

4:30 PM-6:35 PM

Room: Emerald Ballroom - 2nd Floor

Chair: To Be Determined

4:30-4:50 Dimension-independent Sparse Fourier Transform

Michael Kapralov, IBM T.J. Watson Research Center, U.S.; *Ameya Velingker*, Google, Inc., U.S.; Amir Zandieh, École Polytechnique Fédérale de Lausanne, Switzerland

4:55-5:15 Adaptive Sparse Recovery with Limited Adaptivity

Akshay Kamath and Eric Price, University of Texas at Austin, U.S.

5:20-5:40 Efficient Algorithms and Lower Bounds for Robust Linear Regression

Weihao Kong, Stanford University, U.S.; Ilias Diakonikolas, University of Southern California, U.S.; Alistair Stewart, Stanford University, U.S.

5:45-6:05 High-Dimensional Robust Mean Estimation in Nearly-linear Time

Yu Cheng, Duke University, U.S.; Ilias Diakonikolas, University of Southern California, U.S.; Rong Ge, Duke University, U.S.

6:10-6:30 Relative Error Tensor Low Rank Approximation

Zhao Song, University of Texas at Austin, U.S.; David Woodruff, Carnegie Mellon University, U.S.; *Peilin Zhong*, Columbia University, U.S.

Wednesday, January 9

CP46

SODA Session 12B

4:30 PM-6:35 PM

Room: Topaz - 2nd Floor

Chair: To Be Determined

4:30-4:50 Popular Matchings and Limits to Tractability

Yuri Faenza, Columbia University, U.S.; Telikepalli Kavitha, Tata Institute of Fundamental Research, India; Vladlena Powers and Xingyu Zhang, Columbia University, U.S.

4:30-4:50 Popular Matching in Roommates Setting is Np-Hard

Pranabendu Misra and Sushmita Gupta, University of Bergen, Norway; Saket Saurabh, Institute of Mathematical Sciences, India and University of Bergen, Norway; Meirav Zehavi, Ben-Gurion University, Israel

4:55-5:15 A (1+1/e)-Approximation Algorithm for Maximum Stable Matching with One-sided Ties and Incomplete Lists

Chi-Kit Lam and C. Gregory Plaxton, University of Texas at Austin, U.S.

5:20-5:40 Beating Greedy for Stochastic Bipartite Matching

Buddhima Gamlath, Sagar Kale, and Ola Svensson, École Polytechnique Fédérale de Lausanne, Switzerland

5:45-6:05 Stochastic Matching with Few Queries: New Algorithms and Tools

Soheil Behnezhad and Alireza Farhadi, University of Maryland, U.S.; MohammadTaghi Hajiaghayi, University of Maryland, College Park, U.S.; Nima Reyhani, University of Maryland, U.S.

6:10-6:30 Tight Competitive Ratios of Classic Matching Algorithms in the Fully Online Model

Zhiyi Huang, University of Hong Kong, Hong Kong; Binghui Peng, Tsinghua University, China; *Zhihao Gavin Tang*, University of Hong Kong, Hong Kong; Runzhou Tao, Tsinghua University, China; Xiaowei Wu, City University of Hong Kong, Hong Kong; Yuhao Zhang, University of Hong Kong, Hong Kong

Wednesday, January 9

CP47

SODA Session 12C

4:30 PM-6:35 PM

Room: Diamond 1 - 2nd Floor

Chair: To Be Determined

4:30-4:50 Seth Says: Weak Fréchet Distance is Faster, But Only if it is Continuous and in One Dimension

Tim Ophelders, Michigan State University, U.S.; Kevin Buchin and Bettina Speckmann, Eindhoven University of Technology, Netherlands

4:55-5:15 Fréchet Distance under Translation: Conditional Hardness and an Algorithm via Offline Dynamic Grid Reachability

Karl Bringmann, Marvin Künnemann, and *André Nusser*, Max Planck Institute for Informatics, Germany

5:20-5:40 Approximating (k,l)-Center Clustering for Curves

Kevin Buchin, Eindhoven University of Technology, Netherlands; *Anne Driemel*, Universität Bonn, Germany; Joachim Gudmundsson, University of Sydney, Australia; Michael Horton, New York University, U.S. and University of Sydney, Australia; Irina Kostitsyna, Eindhoven University of Technology, Netherlands; Maarten Loffler, Utrecht University, Netherlands; Martijn Struijs, Eindhoven University of Technology, Netherlands

5:45-6:05 Analysis of Ward's Method

Anna Großwendt, Heiko Röglin, Melanie Schmidt, and *Clemens Rösner*, Universität Bonn, Germany

6:10-6:30 Exact Algorithms and Lower Bounds for Stable Instances of Euclidean K-Means

Zachary Friggstad, *Kamyar Khodamoradi*, and Mohammad Salavatipour, University of Alberta, Canada

Speaker Index

siam®

acm Association for Computing Machinery

ACM-SIAM Symposium on Discrete Algorithms

January 6-9, 2019
Westin San Diego
San Diego, California, USA

ALENEX19

Meeting on
Algorithm Engineering & Experiments

January 7-8, 2019
Westin San Diego
San Diego, California, USA

ANALCO19

Meeting on Analytic Algorithmics and Combinatorics

January 6, 2019
Westin San Diego
San Diego, California, USA

Symposium on Simplicity in Algorithms (SOSA19)

V

W

X

Y

Z

Westin San Diego

LEVEL 2

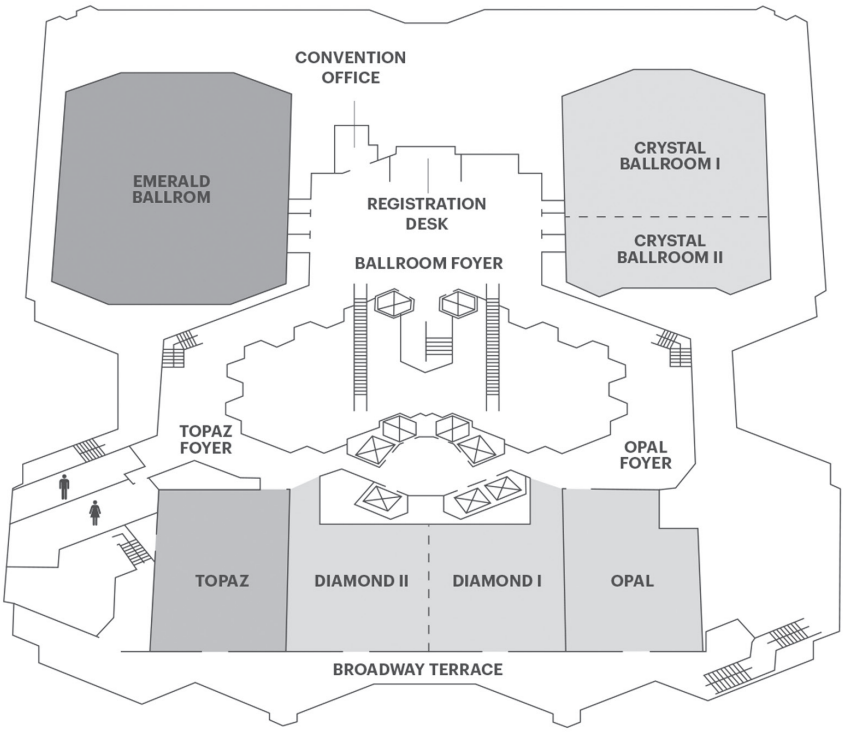

LEVEL 3

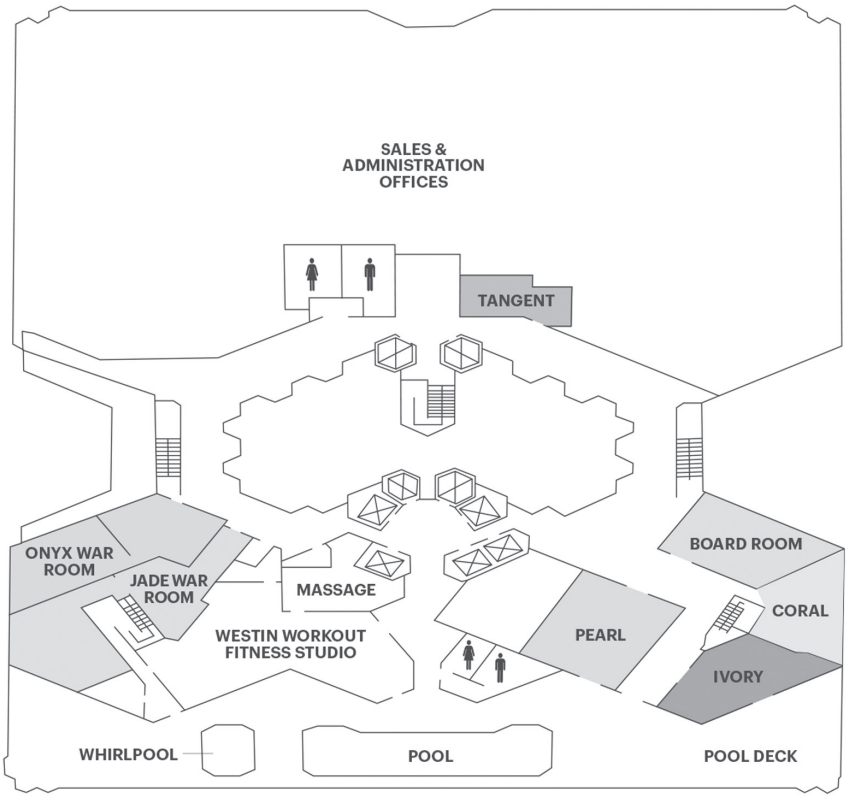

IP1

Towards Analysis of Information Content in Dynamic Networks

Shannon information theory has served as a bedrock for advances in communication and storage systems over the past six decades. However, this theory does not handle well higher order structures (e.g., graphs, geometric structures), temporal aspects (e.g., real-time considerations), or semantics, which are essential aspects of data and information that underlie a broad class of current and emerging data science applications. In this talk, we present some recent results on structural and temporal information in dynamic networks/graphs, in which nodes and edges are added and removed over time. We focus on two related problems: (i) compression of structures – for a given graph model, we exhibit an efficient algorithm for invertibly mapping network structures (i.e., graph isomorphism types) to bit strings of minimum expected length, and (ii) node arrival order inference – for a dynamic graph model, we determine the extent to which the order of node arrivals can be inferred from a snapshot of the graph structure. For both problems, we apply analytic combinatorics, probabilistic, and information-theoretic methods to find statistical limits and efficient algorithms for achieving those limits.

Wojciech Szpankowski
Purdue University
Depratment of Computer Science
spa@cs.purdue.edu

IP2

Recent Advances on the Complexity of Parameterized Counting Problems

Research in the past few decades revealed that a many of the basic algorithmic problems are fixed-parameter tractable with various parameterizations: they can be solved in time $f(k)n^c$, where k is some parameter of the input instance. There have been significant progress both in the design of parameterized algorithms and in understanding the limitations of these techniques. However, up until recent years, there were relatively few algorithmic and complexity results on parameterized counting problems. Similarly to what one experiences in the area of polynomial-time computations, there are natural parameterized problems where finding a single solution is fixed-parameter tractable, but counting the number of solutions is hard. In this talk, I will give a survey of recent developments, including new results that give a clean and unified explanation for the complexity of $\#k$-Path, $\#k$-Matching, and many other parameterized counting problems.

Dániel Marx
Hungarian Academy of Sciences (MTA SZTAKI),
Budapest,Hungary
dmarx@cs.bme.hu

IP3

Inherent Trade-Offs in Algorithmic Fairness

Recent discussion in the public sphere about classification by algorithms has involved tension between competing notions of what it means for such a classification to be fair to different groups. We consider several of the key fairness conditions that lie at the heart of these debates, and discuss recent research establishing inherent trade-offs between these conditions. We also consider a variety of methods for promoting fairness and related notions for classification and selection problems that involve sets rather than just individuals. This talk is based on joint work with Sendhil Mullainathan, Manish Raghavan, and Maithra Raghu.

Jon M. Kleinberg
Cornell University
Dept of Computer Science
kleinber@cs.cornell.edu

IP4

(Hardness of) Approximation and Expansion

Recently, there has been some exciting progress towards proving the unique games conjecture. Is the conjecture true? experts are not sure. Still, it is clearly key to understanding approximation algorithms for many optimization problems, most notably constraint satisfaction. I will describe the background and consequences of the conjecture and then focus on some interesting technical step pertaining to the *expansion* of the so-called Grassman graph. I will then describe a more general phenomenon called high dimensional expansion (HDX). HDX is family of a generalizations of expansion in graphs to hypergraphs. This is a new area of study and has potential for algorithms, as I will try to show.

Irit Dinur
Weizmann Institute
irit.dinur@weizmann.ac.il

CP1

Ranked Schröder Trees

In biology, a phylogenetic tree is a tool to represent the evolutionary relationship between species. Unfortunately, the classical Schröder tree model is not adapted to take into account the chronology between the branching nodes. In particular, it does not answer the question: how many different phylogenetic stories lead to the creation of n species and what is the average time to get there? In this paper, we enrich this model in two distinct ways in order to obtain two ranked tree models for phylogenetics, i.e. models coding chronology. For that purpose, we first develop a model of (strongly) increasing Schröder trees, symbolically described in the classical context of increasing labeling trees. Then we introduce a generalization for the labeling with some unusual order constraint in Analytic Combinatorics (namely the weakly increasing trees). Although these models are direct extensions of the Schröder tree model, it appears that they are also in one-to-one correspondence with several classical combinatorial objects. Through the paper, we present these links, exhibit some parameters in typical large trees and conclude the studies with efficient uniform samplers.

Antoine Genitrini
Sorbonne University / LIP6
antoine.genitrini@lip6.fr

Olivier Bodini
LIPN, Université Paris 13
olivier.bodini@lipn.univ-paris13.fr

Mehdi Naima
Université Paris-Nord

mehdi.naima@lipn.univ-paris13.fr

CP1

Reducing Simply Generated Trees by Iterative Leaf Cutting

We consider a procedure to reduce simply generated trees by iteratively removing all leaves. In the context of this reduction, we study the number of vertices that are deleted after applying this procedure a fixed number of times by using an additive tree parameter model combined with a recursive characterization. Our results include asymptotic formulas for mean and variance of this quantity as well as a central limit theorem.

Benjamin Hackl
Alpen-Adria-Universität Klagenfurt
benjamin.hackl@aau.at

Clemens Heuberger
Alpen-Adria-Universität Klagenfurt, Austria
clemens.heuberger@aau.at

Stephan Wagner
Stellenbosch University, South Africa
swagner@sun.ac.za

CP1

Esthetic Numbers and Lifting Restrictions on the Analysis of Summatory Functions of Regular Sequences

When asymptotically analysing the summatory function of a q-regular sequence in the sense of Allouche and Shallit, the eigenvalues of the sum of matrices of the linear representation of the sequence determine the "shape' (in particular the growth) of the asymptotic formula. Existing general results for determining the precise behavior (including the Fourier coefficients of the appearing fluctuations) have previously been restricted by a technical condition on these eigenvalues. The aim of this work is to lift these restrictions by providing an insightful proof based on generating functions for the main pseudo Tauberian theorem for all cases simultaneously. (This theorem is the key ingredient for overcoming convergence problems in Mellin–Perron summation in the asymptotic analysis.) One example is discussed in more detail: A precise asymptotic formula for the amount of esthetic numbers in the first N natural numbers is presented. Prior to this only the asymptotic amount of these numbers with a given digit-length was known.

Clemens Heuberger, Daniel Krenn
Alpen-Adria-Universität Klagenfurt, Austria
clemens.heuberger@aau.at, math@danielkrenn.at

CP1

Protection Number of Recursive Trees

The protection number of a tree is the minimal distance from its root to a leaf. In this paper we are interested in the protection number of a uniformly chosen random recursive tree of size n. Due to different construction of plane oriented and non-plane recursive trees we consider them separately. We use the singularity analysis of derived generating functions to find out the number of relevant trees, which leads us to the probability distribution of the protection number. Our results are also compared to outcomes of computer simulations.

Zbigniew Golebiewski, Mateusz Klimczak
Wroclaw University of Science and Technology
zbigniew.golebiewski@pwr.edu.pl, mateusz.klimczak@mail.com

CP1

Combinatorics of Nondeterministic Walks of the Dyck and Motzkin Type

This paper introduces *nondeterministic walks*, a new variant of one-dimensional discrete walks. At each step, a nondeterministic walk draws a random set of steps from a predefined set of sets and explores all possible extensions in parallel. We introduce our new model on Dyck steps with the nondeterministic steps $\{-1\}$, $\{1\}$, $\{-1,1\}$ and Motzkin steps with the nondeterministic steps $\{-1\}$, $\{0\}$, $\{1\}$, $\{-1,0\}$, $\{-1,1\}$, $\{0,1\}$, $\{-1,0,1\}$. For general lists of step sets and a given length, we express the generating function of nondeterministic walks where at least one of the walks explored in parallel is a bridge (ends at the origin). In the particular cases of Dyck and Motzkin steps, we also compute the asymptotic probability that at least one of those parallel walks is a meander (stays nonnegative) or an excursion (stays nonnegative and ends at the origin). This research is motivated by the study of networks involving encapsulations and decapsulations of protocols. Our results are obtained using generating functions and analytic combinatorics.

Elie de Panafieu
Nokia Bell Labs France
depanafieuelie@gmail.com

Mohamed Lamine Lamali, Michael Wallner
LaBRI - Université de Bordeaux, France
mohamed-lamine.lamali@u-bordeaux.fr, wallner.michael@gmail.com

CP2

Fast Modular Subset Sum Using Linear Sketching

Given n positive integers, the Modular Subset Sum problem asks if a subset adds up to a given target t modulo a given integer m. This is a natural generalization of the Subset Sum problem (where $m = +\infty$) with ties to additive combinatorics and cryptography. Recently, in [Bringmann, SODA'17] and [Koiliaris and Xu, SODA'17], efficient algorithms have been developed for the non-modular case, running in near-linear pseudo-polynomial time. For the modular case, however, the best known algorithm by Koiliaris and Xu [Koiliaris and Xu, SODA'17] runs in time $\widetilde{O}(m^{5/4})$. In this paper, we present an algorithm running in time $\widetilde{O}(m)$, which matches a recent conditional lower bound of [Abboud et al.'17] based on the Strong Exponential Time Hypothesis. Interestingly, in contrast to most previous results on Subset Sum, our algorithm does not use the Fast Fourier Transform. Instead, it is able to simulate the "textbook' Dynamic Programming algorithm much faster, using ideas from linear sketching. This is one of the first applications of sketching-based techniques to obtain fast algorithms for combinatorial problems in an offline setting.

Kyriakos Axiotis
Massachusetts Institute of Technology
kaxiotis@mit.edu

Arturs Backurs

Toyota Technological Institute at Chicago
backurs@ttic.edu

Ce Jin
Institute for Interdisciplinary Information Sciences
Tsinghua University
jinc16@mails.tsinghua.edu.cn

Christos Tzamos
University of Wisconsin, Madison
tzamos@wisc.edu

Hongxun Wu
Institute for Interdisciplinary Information Sciences
Tsinghua University
wuhx18@mails.tsinghua.edu.cn

CP2

Seth-Based Lower Bounds for Subset Sum and Bicriteria Path

Subset Sum and k-SAT are two of the most extensively studied problems in computer science, and conjectures about their hardness are among the cornerstones of fine-grained complexity. An important open problem in this area is to base the hardness of one of these problems on the other. Our main result is a tight reduction from k-SAT to Subset Sum on dense instances, proving that Bellman's 1962 pseudo-polynomial $O^*(T)$-time algorithm for Subset Sum on n numbers and target T cannot be improved to time $T^{1-\varepsilon} \cdot 2^{o(n)}$ for any $\varepsilon > 0$, unless the Strong Exponential Time Hypothesis (SETH) fails. As a corollary, we prove a "Direct-OR" theorem for Subset Sum under SETH, offering a new tool for proving conditional lower bounds: It is now possible to assume that deciding whether one out of N given instances of Subset Sum is a YES instance requires time $(NT)^{1-o(1)}$. As an application of this corollary, we prove a tight SETH-based lower bound for the classical Bicriteria Path problem, which is extensively studied in Operations Research. We separate its complexity from that of Subset Sum: On graphs with m edges and edge lengths bounded by L, we show that the $O(Lm)$ pseudo-polynomial time algorithm by Joksch from 1966 cannot be improved to $\tilde{O}(L+m)$, in contrast to a recent improvement for Subset Sum (Bringmann, SODA 2017).

Amir Abboud
IBM Research Almaden
abboud@cs.stanford.edu

Karl Bringmann
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
kbringma@mpi-inf.mpg.de

Danny Hermelin, Dvir Shabtay
Ben-Gurion University
hermelin@bgu.ac.il, dvirs@bgu.ac.il

CP2

An Equivalence Class for Orthogonal Vectors

Abstract not available.

Lijie Chen
MIT
lijieche@mit.edu

Ryan Williams
Massachusetts Institute of Technology
rrw@mit.edu

CP2

Fine-grained Complexity Meets IP = PSPACE

In this paper we study the fine-grained complexity of finding exact and approximate solutions to problems in P. Our main contribution is showing reductions from an exact to an approximate solution for a host of such problems. As one (notable) example, we show that the Closest-LCS-Pair problem (Given two sets of strings A and B, compute exactly the maximum LCS(a,b) with $(a, b) \in A \times B$) is equivalent to its constant approximation version under near-linear time reductions. More generally, we identify a class of problems, which we call BP-SAT-Equivalence-Class, comprising both exact and approximate solutions, and show that they are all equivalent under near-linear time reductions. Exploring this class and its properties, we also show: 1. Under the NC-SETH assumption (a significantly more relaxed assumption than SETH), solving any of the problems in this class requires essentially quadratic time. 2. Modest improvements on the running time of known algorithms (shaving log factors) would imply that NEXP is not in non-uniform NC^1. 3. Finally, we leverage our techniques to show new barriers for deterministic approximation algorithms for LCS. At the heart of these new results is a deep connection between interactive proof systems for bounded-space computations and the fine-grained complexity of exact and approximate solutions to problems in P. In particular, our results build on the proof techniques from the classical IP = PSPACE result.

Lijie Chen, Shafi Goldwasser
MIT
lijieche@mit.edu, shafi@theory.csail.mit.edu

Kaifeng Lyu
Tsinghua University
lkf15@mails.tsinghua.edu.cn

Guy Rothblum
Microsoft Research
rothblum@alum.mit.edu

Aviad Rubinstein
Harvard University
aviad@cs.stanford.edu

CP2

A Subquadratic Approximation Scheme for Partition

The subject of this paper is the time complexity of approximating Knapsack, Subset Sum, Partition, and some other related problems. The main result is an $O(n + 1/\varepsilon^{5/3})$ time randomized FPTAS for Partition, which is derived from a certain relaxed form of a randomized FPTAS for Subset Sum. To the best of our knowledge, this is the first NP-hard problem that has been shown to admit a subquadratic time approximation scheme, i.e., one with time complexity of $O((n + 1/\varepsilon)^{2-\delta})$ for some $\delta > 0$. To put these developments in context, note that a quadratic FPTAS for Partition has been known for 40 years. Our main contribution lies in designing a mechanism that reduces an instance of Subset Sum to several simpler instances, each with some special structure, and keeps track of interactions

between them. This allows us to combine techniques from approximation algorithms, pseudo-polynomial algorithms, and additive combinatorics. We also prove several related results. Notably, we improve approximation schemes for 3SUM, (min,+)-convolution, and Tree Sparsity. Finally, we argue why breaking the quadratic barrier for approximate Knapsack is unlikely by giving an $\Omega((n+1/\varepsilon)^{2-o(1)})$ conditional lower bound.

Marcin Mucha
University of Warsaw
Faculty of Mathematics, Informatics and Mechanics
mucha@mimuw.edu.pl

Karol Wegrzycki, Michal Wlodarczyk
University of Warsaw
k.wegrzycki@mimuw.edu.pl,
m.wlodarczyk@mimuw.edu.pl

CP3

A Nearly-Linear Bound for Chasing Nested Convex Bodies

Friedman and Linial introduced the convex body chasing problem to explore the interplay between geometry and competitive ratio in metrical task systems. In convex body chasing, at each time step $t \in \mathbb{N}$, the online algorithm receives a request in the form of a convex body $K_t \subset \mathbb{R}^d$ and must output a point $x_t \in K_t$. The goal is to minimize the total movement between consecutive output points, where the distance is measured in some given norm. This problem is still far from being understood. Recently Bansal et al. gave an $6^d(d!)^2$-competitive algorithm for the *nested* version, where each convex body is contained within the previous one. We propose a different strategy which is $O(d \log d)$-competitive algorithm for this nested convex body chasing problem. Our algorithm works for any norm. This result is almost tight, given an $\Omega(d)$ lower bound for the ℓ_∞ norm.

C.J. Argue
Carnegie Mellon University
Carnegie Mellon University
cargue@andrew.cmu.edu

Sebastien Bubeck
Microsoft Research
sebubeck@microsoft.com

Michael B. Cohen
MIT
micohen@mit.edu

Anupam Gupta
Carnegie Mellon University
anupamg@cs.cmu.edu

Yin Tat Lee
University of Washington
yintat@uw.edu

CP3

Metrical Task Systems on Trees Via Mirror Descent and Unfair Gluing

Abstract not available

Sebastien Bubeck
Microsoft Research
sebubeck@microsoft.com

Michael B. Cohen
MIT
micohen@mit.edu

James R. Lee, Yin Tat Lee
University of Washington
jrl@cs.washington.edu, yintat@uw.edu

CP3

k-Servers with a Smile: Online Algorithms via Projections

We consider the k-server problem on trees and HSTs. We give finite algorithms based on the standard convex-programming primitive of Bregman projections. These algorithms have competitive ratios that match some of the recent results given by Bubeck et al. (STOC 2018), whose algorithms were not finite, but were based using mirror-descent-based continuous dynamics prescribed via a differential inclusion.

Marco Molinaro
Georgia Institute of Technology
School of Industrial and System Engineering
mmolinaro@inf.puc-rio.br

Niv Buchbinder
Statistics and Operations Research Dept.
Tel Aviv University, Israel
niv.buchbinder@gmail.com

Anupam Gupta
Carnegie Mellon University
anupamg@cs.cmu.edu

Joseph Naor
Technion
naor@cs.technion.ac.il

CP3

Elastic Caching

Motivated by applications in cloud computing, we study the classical online caching problem for a cache of variable size, where the algorithm pays a maintenance cost that monotonically increases with cache size. This captures not only the classical setting of a fixed cache size, which corresponds to a maintenance cost of 0 for a cache of size at most k and ∞ otherwise, but also other natural settings in the context of cloud computing such as a concave rental cost on cache size. We call this the *elastic caching* problem. For this problem, we provide several results:

- we give a randomized algorithm with a competitive ratio of $O(\log n)$;
- for concave, or more generally submodular maintenance costs, we give a deterministic algorithm with a competitive ratio of 2;
- when the cost function is any monotone non-negative set function, we give a deterministic n-competitive online algorithm;
- for the offline version of the problem, we give a randomized constant-factor approximation algorithm.

Our algorithms are based on a configuration LP formulation of the problem, for which our main technical contribution is to maintain online a feasible fractional solution that can be converted to an integer solution using existing rounding techniques.

Debmalya Panigrahi
Duke University
debmalya@cs.duke.edu

Anupam Gupta
Carnegie Mellon University
anupamg@cs.cmu.edu

Ravishankar Krishnaswamy
Microsoft Research India
rakri@microsoft.com

Amit Kumar
IIT Delhi
amitk@cse.iitd.ac.in

CP3

A ϕ-Competitive Algorithm for Scheduling Packets with Deadlines

In the *online packet scheduling problem with deadlines* (PacketScheduling, for short), the goal is to schedule transmissions of packets that arrive over time in a network switch and need to be sent across a link. Each packet has a deadline, representing its urgency, and a non-negative weight, that represents its priority. Only one packet can be transmitted in any time slot, so, if the system is overloaded, some packets will inevitably miss their deadlines and be dropped. In this scenario, the natural objective is to compute a transmission schedule that maximizes the total weight of packets which are successfully transmitted. The problem is inherently online, with the scheduling decisions made without the knowledge of future packet arrivals. The central problem concerning PacketScheduling, that has been a subject of intensive study since 2001, is to determine the optimal competitive ratio of online algorithms, namely the worst-case ratio between the optimum total weight of a schedule (computed by an offline algorithm) and the weight of a schedule computed by a (deterministic) online algorithm. We solve this open problem by presenting a ϕ-competitive online algorithm for PacketScheduling (where $\phi \approx 1.618$ is the golden ratio), matching the previously established lower bound.

Pavel Vesely
Department of Computer Science, University of Warwick, UK
vesely@iuuk.mff.cuni.cz

Marek Chrobak
Department of Computer Science
University of California at Riverside, US
marek@cs.ucr.edu

Lukasz Jez
Tel Aviv University
University of Wroclaw, Institute of Computer Science
lje@cs.uni.wroc.pl

Jiri Sgall
Computer Science Institute
Charles University, Czech Republic
sgall@iuuk.mff.cuni.cz

CP4

Multi-unit Supply-monotone Auctions with Bayesian Valuations

We design multi-unit auctions for budget-constrained bidders in the Bayesian setting. This extends the work of Dobzinski, Lavi, and Nisan (Games and Economic Behavior, 2012) and Goel, Mirrokni, and Paes Leme (Games and Economic Behavior, 2015) from the adversarial to the Bayesian setting. Our auctions are supply-monotone, which allows the auction to be run online without knowing the number of items in advance, and achieve asymptotic revenue-optimality compared to the offline revenue-maximizing auction (Pai and Vohra, Journal of Economic Theory, 2014). We also give an efficient algorithm for implementing our auction by using a succinct and efficiently implementable characterization of supply-monotonicity in the Bayesian setting. This provides a generalization of succinct characterizations for single unit allocations due to Border (Journal of Econometric Society, 1991; Economic Theory 2007) and Cai, Daskalakis, and Weinberg (STOC, 2012) to the multi-unit setting.

Yuan Deng, Debmalya Panigrahi
Duke University
ericdy@cs.duke.edu, debmalya@cs.duke.edu

CP4

Tight Revenue Gaps among Simple Mechanisms

We consider a fundamental problem in microeconomics: Selling a single item among a number of buyers whose values are drawn from known independent and regular distributions. There are four widely used and studied mechanisms in this literature: Anonymous Posted-Pricing (AP), Second-Price Auction with Anonymous Reserve (AR), Sequential Posted-Pricing (SPM) and Myerson Auction (OPT). Myerson Auction is optimal but complicated, which also suffers a few issues in practice such as fairness; AP is the simplest mechanism, but its revenue is also the lowest among these four; AR and SPM are of intermediate complexity and revenue. We study the revenue gaps among these four mechanisms, which is defined as the largest ratio between revenues from two mechanisms. We establish two tight ratios and one tighter bound:

- SPM/AP. This ratio studies the power of discrimination in pricing schemes. We obtain the tight ratio of roughly 2.62, closing the previous known bounds $[e/(e-1), e]$.
- AR/AP. This ratio studies the relative power of auction vs. pricing schemes, when no discrimination is allowed. We get the tight ratio of $\pi^2/6 \approx 1.64$, closing the previous known bounds $[e/(e-1), e]$.
- OPT/AR. This ratio studies the power of discrimination in auctions. Previously, the revenue gap is known to be in interval $[2, e]$, and the lower-bound of 2 is conjectured to be tight. We disprove this conjecture by obtaining a better lower-bound of 2.15.

Yaonan Jin
The Hong Kong University of Science and Technology
yjinan@connect.ust.hk

Pinyan Lu
Shanghai University of Finance and Economics
lu.pinyan@mail.shufe.edu.cn

Zhihao Gavin Tang
The University of Hong Kong
zhtang@cs.hku.hk

Tao Xiao
Shanghai Jiao Tong University
xt_1992@sjtu.edu.cn

CP4

Assignment Mechanisms under Distributional Constraints

We study the assignment problem of objects to agents with heterogeneous preferences under distributional constraints. Each agent is associated with a publicly known type and has a private ordinal ranking over objects. We are interested in assigning as many agents as possible. Our first contribution is a generalization of the well-known and widely used serial dictatorship. Our mechanism maintains several desirable properties of serial dictatorship, including strategyproofness, Pareto efficiency, and computational tractability while satisfying the distributional constraints with a small error. We also propose a generalization of the probabilistic serial algorithm, which finds an ordinally efficient and envy-free assignment, and also satisfies the distributional constraints with a small error. We show, however, that no ordinally efficient and envy-free mechanism is also weakly strategyproof. Both of our algorithms assign at least the same number of students as the optimum fractional assignment.

Itai Ashlagi
Stanford University
iashlagi@stanford.edu

Amin Saberi
Management Science and Engineering
Stanford University
saberi@stanford.edu

Ali Shameli
Stanford University
shameli@stanford.edu

CP4

Correlation-robust Analysis of Single Item Auction

We investigate the problem of revenue maximization in single-item auction within the new correlation-robust framework proposed by Carroll [2017] and further developed by Gravin and Lu [2018]. In this framework the auctioneer is assumed to have only partial information about marginal distributions, but does not know the dependency structure of the joint distribution. The auctioneer's revenue is evaluated in the worst-case over the uncertainty of possible joint distribution. For the problem of optimal auction design in the correlation robust-framework we observe that in most cases the optimal auction does not admit a simple form like the celebrated Myerson's auction for independent valuations. We analyze and compare performances of several DSIC mechanisms used in practice. Our main set of results concern the sequential posted-price mechanism (SPM). We show that SPM achieves a constant (4.78) approximation to the optimal correlation-robust mechanism. We also show that in the symmetric (anonymous) case when all bidders have the same marginal distribution, (i) SPM has almost matching worst-correlation revenue as any second price auction with common reserve price, and (ii) when the number of bidders is large, SPM converges to optimum. In addition, we extend some results on approximation and computational tractability for lookahead auctions to the correlation-robust framework.

Xiaohui Bei
Nanyang Technological University
xhbei@ntu.edu.sg

Nick Gravin, Pinyan Lu
Shanghai University of Finance and Economics
nikolai@mail.shufe.edu.cn, lu.pinyan@mail.shufe.edu.cn

Zhihao Gavin Tang
The University of Hong Kong
zhtang@cs.hku.hk

CP4

Dynamic Double Auctions: Towards First Best

We study the problem of designing dynamic double auctions for two-sided markets in which a platform intermediates the trade between one seller offering independent items to multiple buyers, repeatedly over a finite horizon. Motivated by online advertising and ride-hailing markets, we seek to design mechanisms satisfying the following properties: no positive transfers, i.e., the platform never asks the seller to make payments nor buyers are ever paid and periodic individual rationality, i.e., every agent should derive a non-negative utility from every trade opportunity. We provide simple mechanisms satisfying these requirements that are asymptotically efficient and budget-balanced with high probability as the number of trading opportunities grows. Moreover, we show that the average expected profit obtained by the platform under these mechanisms asymptotically approaches first-best (the maximum possible welfare generated by the market) either in the case of one buyer and one seller, or multiple buyers and one seller with deterministic values.

Santiago Balseiro
Columbia University
srb2155@columbia.edu

Vahab Mirrokni
Google Inc.
mirrokni@google.com

Renato Paes Leme, Song Zuo
Google Research
renatoppl@google.com, szuo@google.com

CP5

QuickSort: Improved Right-Tail Asymptotics for the Limiting Distribution, and Large Deviations

We refine asymptotic logarithmic upper bounds – extending from one term to three – produced by Svante Janson (2015) on the right tail of the limiting QuickSort distribution function F and by Fill and Hung (2018) on the right tails of the corresponding density f and of the absolute derivatives of f of each order. All our results match two-term lower bounds for the functions in question to two terms and match conjectured asymptotic expansions to three terms. Using the refined asymptotic bounds on F, we derive right-tail large deviation (LD) results for the distribution of the number of comparisons required by Quick-

Sort that sharpen somewhat the two-sided LD results of McDiarmid and Hayward (1996).

James Allen Fill, Wei-Chun Hung
Johns Hopkins University
jimfill@jhu.edu, whung6@jhu.edu

CP5

Moments of Select Sets

We analyze a selection procedure introduced by Krieger, Pollak, and Samuel–Cahn. It retains an item if it is among the top $100p$ percent, as compared to the items that have been accepted so far. Gaither and Ward analyzed the average behavior of the number of items selected. We present the asymptotic properties of the higher moments of the number of items retained by the selection procedure. We derive a general formula for the moments. To demonstrate the complexity of these moments, we present the exact first-order asymptotic growth of some of these moments, for various rational values of p.

Simon H. Langowski, Mark Daniel Ward
Purdue University
Department of Statistics
slangows@purdue.edu, mdw@purdue.edu

CP5

Median-of-K Jumplists and Dangling-Min BSTs

We extend randomized jumplists introduced by Brönnimann et al. (STACS 2003) to choose jump-pointer targets as median of a small sample for better search costs, and present randomized algorithms with expected $O(\log n)$ time complexity that maintain the probability distribution of jump pointers upon insertions and deletions. We analyze the expected costs to search, insert and delete a random element, and we show that omitting jump pointers in small sublists hardly affects search costs, but significantly reduces the memory consumption. We use a bijection between jumplists and "dangling-min BSTs', a variant of (fringe-balanced) binary search trees for the analysis. Despite their similarities, some standard analysis techniques for search trees fail for dangling-min trees (and hence for jumplists).

Sebastian Wild
David R. Cheriton School of Computer Science
University of Waterloo
wild@uwaterloo.ca

Markus Nebel
Universität Bielefeld
nebel@techfak.uni-bielefeld.de

Elisabeth Neumann
Carl-Friedrich-Gauß-Fakultät
Technische Universität Braunschweig, Germany
e.neumann@tu-braunschweig.de

CP5

Sesquickselect: One and a Half Pivots for Cache-Efficient Selection

Because of unmatched improvements in CPU performance, memory transfers have become a bottleneck of program execution. As discovered in recent years, this also affects sorting in internal memory. Since partitioning around several pivots reduces overall memory transfers, we have seen renewed interest in multiway Quicksort. Here, we analyze in how far multiway partitioning helps in Quickselect. We compute the expected number of comparisons and scanned elements (approximating memory transfers) for a generic class of (non-adaptive) multiway Quickselect and show that three or more pivots are not helpful, but two pivots are. Moreover, we consider "adaptive' variants which choose partitioning and pivot-selection methods in each recursive step from a finite set of alternatives depending on the current (relative) sought rank. We show that "Sesquickselect', a new Quickselect variant that uses either one or two pivots, makes better use of small samples w.r.t. memory transfers than other Quickselect variants.

Sebastian Wild
David R. Cheriton School of Computer Science
University of Waterloo
wild@uwaterloo.ca

Conrado Martínez
Department of Computer Science
Univ. Politècnica de Catalunya
conrado@cs.upc.edu

Markus Nebel
Universität Bielefeld
nebel@techfak.uni-bielefeld.de

CP6

Stochastic Submodular Cover with Limited Adaptivity

In the submodular cover problem, we are given a non-negative monotone submodular function f over a ground set E of items, and the goal is to choose a smallest subset $S \subseteq E$ such that $f(S) = Q$ where $Q = f(E)$. In the stochastic version of the problem, we are given m stochastic items which are different random variables that independently realize to some item in E, and the goal is to find a smallest set of stochastic items whose realization R satisfies $f(R) = Q$. The problem captures as a special case the stochastic set cover problem and more generally, stochastic covering integer programs. We study the role of adaptivity in the stochastic submodular cover problem. We define an r-round adaptive algorithm to be an algorithm that chooses a permutation of all available items in each round $k \in [r]$, and a threshold τ_k, and realizes items in the order specified by the permutation until the function value exceeds τ_k. Our main result is that for any integer r, there exists a poly-time r-round adaptive algorithm for stochastic submodular cover whose expected cost is $\tilde{O}(Q^{1/r})$ times the expected cost of a fully adaptive algorithm. We complement this result by showing that for any r, there exist instances of the stochastic submodular cover problem where no r-round adaptive algorithm can achieve better than $\Omega(Q^{1/r})$ approximation to the expected cost of a fully adaptive algorithm.

Arpit Agarwal
PhD Student
University of Pennsylvania
aarpit@cis.upenn.edu

Sepehr Assadi, Sanjeev Khanna
University of Pennsylvania

sassadi@cis.upenn.edu, sanjeev@cis.upenn.edu

CP6

An Exponential Speedup in Parallel Running Time for Submodular Maximization Without Loss in Approximation

Abstract not available

Eric Balkanski, Aviad Rubinstein, Yaron Singer
Harvard University
ericbalkanski@g.harvard.edu, aviad@cs.stanford.edu, yaron@seas.harvard.edu

CP6

Submodular Maximization with Nearly-optimal Approximation and Adaptivity in Nearly-linear Time

In this paper, we study the tradeoff between the approximation guarantee and adaptivity for the problem of maximizing a monotone submodular function subject to a cardinality constraint. The adaptivity of an algorithm is the number of sequential rounds of queries it makes to the evaluation oracle of the function, where in every round the algorithm is allowed to make polynomially-many parallel queries. Adaptivity is an important consideration in settings where the objective function is estimated using samples and in applications where adaptivity is the main running time bottleneck. Previous algorithms achieving a nearly-optimal $1 - 1/e - \epsilon$ approximation require $\Omega(n)$ rounds of adaptivity. In this work, we give the first algorithm that achieves a $1 - 1/e - \epsilon$ approximation using $O(\ln n/\epsilon^2)$ rounds of adaptivity. The number of function evaluations and additional running time of the algorithm are $O(n \operatorname{poly}(\log n, 1/\epsilon))$.

Alina Ene
Boston University
aene@bu.edu

Huy Nguyen
Northeastern University
hlnguyen@cs.princeton.edu

CP6

Submodular Maximization with Nearly Optimal Approximation, Adaptivity and Query Complexity

Submodular optimization generalizes many classic problems in combinatorial optimization and has recently found a wide range of applications in machine learning (e.g., feature engineering and active learning). For many large-scale optimization problems, we are often concerned with the adaptivity complexity of an algorithm, which quantifies the number of sequential rounds where polynomially-many independent function evaluations can be executed in parallel. While low adaptivity is ideal, it is not sufficient for a distributed algorithm to be efficient, since in many practical applications of submodular optimization the number of function evaluations becomes prohibitively expensive. Motivated by these applications, we study the adaptivity and query complexity of adaptive submodular optimization. Our main result is a distributed algorithm for maximizing a monotone submodular function with cardinality constraint k that achieves a $(1 - 1/e - \varepsilon)$-approximation in expectation. This algorithm runs in $O(\log(n))$ adaptive rounds and makes $O(n)$ calls to the function evaluation oracle in expectation. The approximation guarantee and query complexity are optimal, and the adaptivity is nearly optimal. Moreover, the number of queries is substantially less than in previous works. Last, we extend our results to the submodular cover problem to demonstrate the generality of our algorithm and techniques.

Matthew Fahrbach
Georgia Institute of Technology
matthew.fahrbach@gatech.edu

Vahab Mirrokni
Google Inc.
mirrokni@google.com

Morteza Zadimoghaddam
Google
zadim@google.com

CP6

Deterministic (1/2 + ϵ)-Approximation for Submodular Maximization over a Matroid

In this talk, we will consider the problem of maximizing a monotone submodular function subject to a matroid constraint and present a deterministic algorithm that achieves $(1/2 + \epsilon)$-approximation for the problem. This algorithm is the first deterministic algorithm known to improve over the 1/2-approximation ratio of the classical greedy algorithm proved by Nemhauser, Wolsely and Fisher in 1978.

Niv Buchbinder
Statistics and Operations Research Dept.
Tel Aviv University, Israel
niv.buchbinder@gmail.com

Moran Feldman, Mohit Garg
Open University of Israel
moranfe@openu.ac.il, bootup.c@gmail.com

CP6

Stochastic ℓ_p Load Balancing and Moment Problems via the L-Function Method

This paper considers stochastic optimization problems whose objective functions involve powers of random variables. For example, consider the classic Stochastic lp Load Balancing Problem (SLBp): There are m machines and n jobs, and known independent random variables Y_{ij} decribe the load incurred on machine i if we assign job j to it. The goal is to assign each jobs to machines in order to minimize the expected l_p-norm of the total load on the machines. While convex relaxations represent one of the most powerful algorithmic tools, in problems such as SLBp the main difficulty is to capture the objective function in a way that only depends on each random variable separately. We show how to capture p-power-type objectives in such separable way by using the L-function method, introduced by Latała to relate the moment of sums of random variables to the individual marginals. We show how this quickly leads to a constant-factor approximation for very general subset selection problem with p-moment objective. Moreover, we give a constant-factor approximation for SLBp, improving on the recent $O(p/\ln p)$-approximation of [Gupta et al., SODA 18]. Here the application of the method is much more involved. In particular, we need to sharply connect the expected l_p-norm of a random vector with the p-moments of its marginals (machine loads), taking into ac-

count simultaneously the different scales of the loads that are incurred by an unknown assignment.

Marco Molinaro
Georgia Institute of Technology
School of Industrial and System Engineering
mmolinaro@inf.puc-rio.br

CP6

Submodular Function Maximization in Parallel via the Multilinear Extension

Balkanski and Singer [5] recently initiated the study of adaptivity (or parallelism) for constrained submodular function maximization, and studied the setting of a cardinality constraint. Very recent improvements for this problem by Balkanski, Rubinstein, and Singer [6] and Ene and Nguyen [21] resulted in a near-optimal $(1 - 1/e - \epsilon)$-approximation in $O(\log n/\epsilon^2)$ rounds of adaptivity. Partly motivated by the goal of extending these results to more general constraints, we describe parallel algorithms for approximately maximizing the multilinear relaxation of a monotone submodular function subject to packing constraints. Formally our problem is to maximize $F(x)$ over $x \in [0,1]^n$ subject to $Ax \leq 1$ where F is the multilinear relaxation of a monotone submodular function. Our algorithm achieves a near-optimal $(1 - 1/e - \epsilon)$-approximation in $O(\log^2 m \log n/\epsilon^4)$ rounds where n is the cardinality of the ground set and m is the number of packing constraints. For many constraints of interest, the resulting fractional solution can be rounded via known randomized rounding schemes that are oblivious to the specific submodular function. We thus derive randomized algorithms with polylogarithmic adaptivity for a number of constraints including partition and laminar matroids, matchings, knapsack constraints, and their intersections.

Chandra Chekuri, Kent Quanrud
University of Illinois at Urbana-Champaign
chekuri@illinois.edu, quanrud2@illinois.edu

CP7

Approximate Nearest Neighbor Searching with Non-euclidean and Weighted Distances

We present a new approach to ε-approximate nearest-neighbor searching in R^d for fixed d under a variety of distance functions including the Bregman divergence, the Mahalanobis distance, the Minkowski metric, multiplicative weights, and convex scaling distance functions. The most efficient data structure previously known for the Euclidean metric resisted generalization because of its reliance on the lifting transformation, and the best known approaches for other distance functions are far less efficient. We circumvent the reliance on the lifting transformation by a careful application of *convexification*. While this technique is well known in the field of non-linear optimization, it appears to be relatively new to computational geometry. Under mild assumptions on the distance functions in question, the proposed data structures answer queries in logarithmic time with storage $O(n \log(1/\varepsilon)/\varepsilon^{d/2})$, where n is the number of data points. This nearly matches the best known results for the Euclidean metric.

Ahmed Abdelkader
University of Maryland
akader@cs.umd.edu

Sunil Arya
The Hong Kong University of Science and Technology
arya@cse.ust.hk

Guilherme D. da Fonseca
Université Clermont Auvergne, LIMOS
France
fonseca@isima.fr

David M. Mount
University of Maryland
mount@cs.umd.edu

CP7

New Lower Bounds for the Number of Pseudoline Arrangements

Arrangements of lines and pseudolines are fundamental objects in discrete and computational geometry. They also appear in other areas of computer science, such as the study of sorting networks. Let B_n be the number of arrangements of n pseudolines and let $b_n = \log_2 B_n$. The problem of estimating B_n was posed by Knuth in 1992. Knuth conjectured that $b_n \leq \binom{n}{2} + o(n^2)$ and also derived the first upper and lower bounds: $b_n \leq 0.7924(n^2 + n)$ and $b_n \geq n^2/6 - O(n)$. The upper bound underwent several improvements, $b_n \leq 0.6988n^2$ (Felsner, 1997), and $b_n \leq 0.6571n^2$ (Felsner and Valtr, 2011), for large n. Here we show that $b_n \geq cn^2 - O(n \log n)$ for some constant $c > 0.2053$. In particular, $b_n \geq 0.2053\, n^2$ for large n. This improves the previous best lower bound, $b_n \geq 0.1887n^2$, due to Felsner and Valtr (2011). Our arguments are elementary and geometric in nature. Further, our constructions are likely to spur new developments and improved lower bounds for related problems, such as in topological graph drawings.

Adrian Dumitrescu
Department of Computer Science
University of Wisconsin-Milwaukee
dumitres@uwm.edu

Ritankar Mandal
University of Wisconsin–Milwaukee, USA
rmandal@uwm.edu

CP7

Optimal Algorithm for Geodesic Nearest-point Voronoi Diagrams in Simple Polygons

Given a set of m point sites in a simple polygon, the geodesic nearest-point Voronoi diagram of the sites partitions the polygon into m Voronoi cells, one cell per site, such that every point in a cell has the same nearest site under the geodesic metric. In this paper, we present an $O(n + m \log m)$-time algorithm for computing the geodesic nearest-point Voronoi diagram of m points in a simple n-gon. This matches the best known lower bound of $\Omega(n + m \log m)$ as well as improving the previously best known algorithms which take time $O(n + m \log m + m \log^2 n)$ and $O(n \log n + m \log m)$. This answers the longstanding question whether the geodesic nearest-point Voronoi diagram can be computed optimally, which was explicitly posed by Aronov [Algorithmica, 1989] and Mitchell [Handbook of Computational Geometry, 2000].

Eunjin Oh
Max Planck Institute for Informatics

eoh@mpi-inf.mpg.de

CP7

Extremal and Probabilistic Results for Order Types

A *configuration* is a finite set of points in the plane. Two configurations A and B have the same *order type* if there exists a bijection between them preserving the orientation of every ordered triple. We investigate extremal and probabilistic problems related to configurations in general position. We focus on problems involving forbidden configurations or monotone/hereditary properties. Given a configuration B, we consider the property of being "B-free': a configuration A is *B-free* if no subset of points of A has the same order type as B. We prove a significant bound on the number of B-free N-point configurations contained in the $m \times m$ grid $[m]^2$ for arbitrary configurations B. We consider random N-point configurations U in the unit square, in which each of the N points is chosen uniformly at random and independently of all other points. The above-mentioned enumeration result for B-free configurations in the grid is then used to prove strong bounds for the probability that the random set U should be B-free for any given B. We also investigate the threshold function $N_0 = N_0(n)$ for the property that U should be *n-universal*, that is, should contain all n-point configurations in general position. As it turns out, $N_0 = N_0(n)$ is doubly exponential in n. Our arguments are mostly geometric and combinatorial, with the recent container method playing an important role.

Jie Han
University of Rhode Island
jie_han@uri.edu

Yoshiharu Kohayakawa
Instituto de Matematica e Estatistica
Universidade de São Paulo, Brazil
yoshi@ime.usp.br

Marcelo Sales
Emory University
mtsales@emory.edu

Henrique Stagni
University of Sao Paulo
stagni@ime.usp.br

CP7

Colored Range Closest-Pair Problem under General Distance Functions

The range closest-pair (RCP) problem is the range-search version of the classical closest-pair problem, which aims to store a given dataset of points in some data structure such that when a query range X is specified, the closest pair of points contained in X can be reported efficiently. A natural generalization of the RCP problem is the colored RCP (CRCP) problem in which the given data points are colored and the goal is to find the closest bichromatic pair contained in the query range. All the previous work on the RCP problem was restricted to the uncolored version and the Euclidean distance function. In this paper, we make the first progress on the CRCP problem. We investigate the problem under a general distance function induced by a monotone norm; this covers all L_p-metrics for $p > 0$ and the L_∞-metric. We design efficient $(1+\varepsilon)$-approximate CRCP data structures for orthogonal queries in the plane, where $\varepsilon > 0$ is a prespecified parameter. The highlights are two data structures for answering rectangle queries, one of which uses $O(\varepsilon^{-1} n \log^4 n)$ space and $O(\log^4 n + \varepsilon^{-1} \log^3 n + \varepsilon^{-2} \log n)$ query time while the other uses $O(\varepsilon^{-1} n \log^3 n)$ space and $O(\log^5 n + \varepsilon^{-1} \log^4 n + \varepsilon^{-2} \log^2 n)$ query time. In addition, we also apply our techniques to the CRCP problem in higher dimensions, obtaining efficient data structures for slab, 2-box, and 3D dominance queries.

Jie Xue
University of Minnesota, Twin Cities
xuexx193@umn.edu

CP8

An Algorithmic Blend of LPs and Ring Equations for Promise CSPs

Promise CSPs are a relaxation of constraint satisfaction problems where the goal is to find an assignment satisfying a relaxed version of the constraints. Promise CSPs include approximate graph coloring, discrepancy minimization, and interesting variants of satisfiability. Similar to CSPs, the tractability of promise CSPs can be tied to the structure of operations on the solution space called polymorphisms, though in the promise world these operations are much less constrained. In previous work, we classified Boolean promise CSPs when the constraint predicates are symmetric. In this work, we vastly generalize these algorithmic results. Specifically, we show that promise CSPs that admit a family of "regional-periodic" polymorphisms are in P, assuming that determining which region a point is in can be computed in polynomial time. Our algorithm is based on a novel combination of linear programming and solving linear systems over rings. We also abstract a framework based on reducing a promise CSP to a CSP over an infinite domain, solving it there, and then rounding the solution to an assignment for the promise CSP instance. The rounding step is intimately tied to the family of polymorphisms and clarifies the connection between polymorphisms and algorithms in this context. As a key ingredient, we introduce the technique of finding a solution to an LP with integer coefficients that lies in a different ring to bypass ad-hoc adjustments for lying on a rounding boundary.

Joshua Brakensiek, Venkatesan Guruswami
Carnegie Mellon University
jbrakens@stanford.edu, venkatg@cs.cmu.edu

CP8

Perfect Matchings, Rank of Connection Tensors and Graph Homomorphisms

We develop a theory of graph algebras over general fields. This is modeled after the theory developed by Freedman, Lovász and Schrijver in "[Michael Freedman, László Lovász, Alexander Schrijver, Reflection positivity, rank connectivity, and homomorphism of graphs]" for connection matrices, in the study of graph homomorphism functions over real edge weight and positive vertex weight. We introduce connection tensors for graph properties. This notion naturally generalizes the concept of connection matrices. It is shown that counting perfect matchings, and a host of other graph properties naturally defined as Holant problems (edge models), cannot be expressed by graph homomorphism functions over the complex numbers (or even more general fields). Our necessary and sufficient condition in terms of connection tensors is a simple exponential

rank bound. It shows that positive semidefiniteness is not needed in the more general setting.

Jin-Yi Cai
University of Wisconsin, Madison
jyc@cs.wisc.edu

Artsiom Hovarau
University of Wisconsin-Madison
hovarau@cs.wisc.edu

CP8

Probabilistic Tensors and Opportunistic Boolean Matrix Multiplication

We introduce probabilistic extensions of classical deterministic measures of algebraic complexity of a tensor, such as the rank and the border rank. We show that these probabilistic extensions satisfy various natural and algorithmically serendipitous properties, such as submultiplicativity under taking of Kronecker products. Furthermore, the probabilistic extensions enable improvements over their deterministic counterparts for specific tensors of interest, starting from the tensor $\langle 2,2,2\rangle$ that represents 2×2 matrix multiplication. While it is well known that the (deterministic) tensor rank and border rank satisfy

$$\mathrm{rk}\,\langle 2,2,2\rangle = 7 \qquad \text{and} \qquad \underline{\mathrm{rk}}\,\langle 2,2,2\rangle = 7$$

[V. Strassen, *Numer. Math.* 13 (1969); J. E. Hopcroft and L. R. Kerr, *SIAM J. Appl. Math.* 20 (1971); S. Winograd, *Linear Algebra Appl.* 4 (1971); J. M. Landsberg, *J. AMS* 19 (2006)], we show that the *probabilistic* tensor rank and border rank satisfy

$$\widetilde{\mathrm{rk}}\,\langle 2,2,2\rangle \le 6+\frac{6}{7} \qquad \text{and} \qquad \widetilde{\underline{\mathrm{rk}}}\,\langle 2,2,2\rangle \le 6+\frac{2}{3}\,.$$

By submultiplicativity, this leads immediately to novel randomized algorithm designs, such as algorithms for Boolean matrix multiplication as well as detecting and estimating the number of triangles and other subgraphs in graphs.

Petteri Kaski, Matti Karppa
Aalto University
Department of Computer Science
petteri.kaski@aalto.fi, matti.karppa@aalto.fi

CP8

The Complexity of the Ideal Membership Problem for Constrained Problems Over the Boolean Domain

Given an ideal I and a polynomial f the Ideal Membership Problem is to test if $f \in I$. This problem is a fundamental algorithmic problem with important applications and notoriously intractable. We study the complexity of the Ideal Membership Problem for combinatorial ideals that arise from constrained problems over the Boolean domain. As our main result, we identify the precise borderline of tractability. By using Gröbner bases techniques, we generalize Schaefer's dichotomy theorem [STOC, 1978] which classifies all Constraint Satisfaction Problems over the Boolean domain to be either in P or NP-hard. This paper is motivated by the pursuit of understanding the recently raised issue of bit complexity of Sum-of-Squares proofs [O'Donnell, ITCS, 2017]. Raghavendra and Weitz [ICALP, 2017] show how the Ideal Membership Problem tractability for combinatorial ideals implies bounded coefficients in Sum-of-Squares proofs.

Monaldo Mastrolilli
IDSIA, University of Lugano
monaldo@idsia.ch

CP8

Fast Greedy for Linear Matroids

A fundamental algorithmic result for matroids is that the maximum weight base can be computed using the greedy algorithm. For explicitly represented matroids an important question is the time complexity of computing such a base. It is known that one can compute it in time almost linear in the number of non-zero entries of the linear representation plus r^ω, where r is the rank of the matroid and ω is the matrix multiplication exponent. In this work, we give an alternative algorithm for the same task.

Huy L. Nguyen
Northeastern University
huylenguyen@gmail.com

CP9

Subcritical Random Hypergraphs, High-Order Components, and Hypertrees

In the binomial random graph $G(n,p)$, when p changes from $(1-\varepsilon)/n$ (subcritical case) to $1/n$ and then to $(1+\varepsilon)/n$ (supercritical case) for $\varepsilon > 0$, with high probability the order of the largest component increases smoothly from $O(\varepsilon^{-2}\log(\varepsilon^3 n))$ to $\Theta(n^{2/3})$ and then to $(1\pm o(1))2\varepsilon n$. As a natural generalisation of random graphs and connectedness, we consider the binomial random k-uniform hypergraph $H^k(n,p)$ (where each k-tuple of vertices is present as a hyperedge with probability p independently) and the following notion of high-order connectedness. Given an integer $1 \le j \le k-1$, two sets of j vertices are called *j-connected* if there is a walk of hyperedges between them such that any two consecutive hyperedges intersect in at least j vertices. A j-connected component is a maximal collection of pairwise j-connected j-tuples of vertices. Recently, the threshold for the appearance of the giant j-connected component in $H^k(n,p)$ and its order were determined. In this article, we take a closer look at the subcritical random hypergraph. We determine the structure and size (*i.e.* number of hyperedges) of the largest j-connected components, with the help of a certain class of 'hypertrees' and related objects. In our proofs, we combine various probabilistic and enumerative techniques, such as generating functions and couplings with branching processes.

Wenjie Fang, Oliver Cooley, Nicola Del Giudice
Graz University of Technology
fang@math.tugraz.at, cooley@math.tugraz.at,
delgiudice@math.tugraz.at

Mihyun Kang
Technische Universitaet Graz
Institut fuer Optimierung und Diskrete Mathematik
kang@math.tugraz.at

CP9

Degree Distributions of Generalized Hooking Networks

A hooking network is grown from a set of graphs called blocks, each block with a labelled vertex called a hook. At each step in the growth of the network, a vertex called a latch is chosen from the hooking network, and a block is attached by joining the hook of the block with the latch.

These graphs generalize trees, which are hooking networks grown from a single edge as the only block. Using Pólya urns, we show multivariate normal limit laws for the degree distributions of hooking networks. We extend previous results by allowing for more than one block in the growth of the network and by studying arbitrarily large degrees.

Colin Desmarais
Uppsala University
colin.desmarais@math.uu.se

Cecilia Holmgren
Uppsala University
Department of Mathematics
cecilia.holmgren@math.uu.se

CP9

When Does Hillclimbing Fail on Monotone Functions?

Hillclimbing is an essential part of optimization. An important benchmark for hillclimbing algorithms on functions $f : \{0,1\}^n \to \mathbf{R}$ are (strictly) *montone* functions, on which a surprising number of hillclimbers fail to be efficient. For example, the $(1+1)$-Evolutionary Algorithm is a standard hillclimber which flips each bit independently with probability c/n in each round. Perhaps surprisingly, this algorithm shows a phase transition: it optimizes any such monotone function in quasilinear time if $c < 1$, while, on the other hand, it is known how to construct monotone functions for which the algorithm needs exponential time if c is sufficiently large. As the previsouly known argument for quasilinear running time breaks down at $c = 1$, it is natural to suspect that the phase transition should occur at this value. We here present a new runtime analysis for the $(1+1)$-EA which shows that this is not the case. More precisely, there is a $c_0 > 1$ such that the running time of the algorithm remains quasilinear for $c < c_0$. Our proof is based on Moser's entropy compression argument. That is, we show that a long runtime would allow us to encode the random steps of the algorithm with less bits than their entropy.

Anders Martinsson
Department of Computer Science
ETH Zürich
maanders@inf.ethz.ch

CP9

Arithmetic Progression Hypergraphs: Examining the Second Moment Method

In many data structure settings, it has been shown that using "double hashing" in place of standard hashing, by which we mean choosing multiple hash values according to an arithmetic progression instead of choosing each hash value independently, has asymptotically negligible difference in performance. We attempt to extend these ideas beyond data structure settings by considering how threshold arguments based on second moment methods can be generalized to "arithmetic progression" versions of problems. With this motivation, we define a novel "quasi-random" hypergraph model, random arithmetic progression (AP) hypergraphs, which is based on edges that form arithmetic progressions and unifies many previous problems. Our main result is to show that second moment arguments for 3-NAE-SAT and 2-coloring of 3-regular hypergraphs extend to the double hashing setting. We can generalize these results to larger sized hyperedges, when randomly chosen hyperedges satisfy an appropriate notion limited independence. We leave several open problems related to these quasi-random hypergraphs and the thresholds of associated problem variations.

Michael Mitzenmacher
Harvard University
michaelm@eecs.harvard.edu

CP9

Random Walks on Graphs: New Bounds on Hitting, Meeting, Coalescing and Returning

We prove new results on lazy random walks on finite graphs. To start, we obtain new estimates on return probabilities $P^t(x,x)$ and the maximum expected hitting time t_{hit}, both in terms of the relaxation time. We also prove a discrete-time version of the first-named author's "Meeting time lemma" that bounds the probability of a random walk hitting a deterministic trajectory in terms of hitting times of static vertices. The meeting time result is then used to bound the expected full coalescence time of multiple random walks over a graph. This last theorem is a discrete-time version of a result by the first-named author, which had been previously conjectured by Aldous and Fill. Our bounds improve on recent results by Lyons and Oveis-Gharan; Kanade et al; and (in certain regimes) Cooper et al.

Roberto Oliveira
IMPA, Rio de Janeiro
rimfo@impa.br

Yuval Peres
Microsoft Research, Redmond
peres@microsoft.com

CP10

Finding Maximal Sets of Laminar 3-Separators in Planar Graphs in Linear Time

We consider decomposing a 3-connected planar graph G using laminar separators of size three. We show how to find a maximal set of laminar 3-separators in such a graph in linear time. We also discuss how to find maximal laminar set of 3-separators from special families. For example we discuss non-trivial cuts, ie. cuts which split G into two components of size at least two. For any vertex v, we also show how to find a maximal set of 3-separators disjoint from v which are laminar and satisfy: every vertex in a separator X has two neighbours not in the unique component of $G - X$ containing v. In all cases, we show how to construct a corresponding tree decomposition of adhesion three. Our new algorithms form an important component of recent methods for finding disjoint paths in nonplanar graphs.

David Eppstein
University of California, Irvine
Department of Computer Science
eppstein@uci.edu

CP10

A Faster Algorithm for Minimum-Cost Bipartite Matching in Minor-Free Graphs

We give an $\tilde{O}(n^{7/5}\log(nC))$-time algorithm to compute a minimum-cost maximum cardinality matching (optimal

matching) in K_h-minor free graphs with $h = O(1)$ and integer edge weights having magnitude at most C. This improves upon the $\tilde{O}(n^{10/7} \log C)$ algorithm of Cohen et al. [SODA 2017] and the $O(n^{3/2} \log(nC))$ algorithm of Gabow and Tarjan [SIAM J. Comput. 1989]. For a graph with m edges and n vertices, the Hopcroft-Karp [SIAM J. Comput. 1973], and Gabow-Tarjan algorithms compute, in each phase, a maximal set of vertex-disjoint shortest augmenting paths (for appropriately defined costs) in $O(m)$ time. These algorithms arrive at an optimal matching after $O(\sqrt{n})$ phases with a total execution time of $O(m\sqrt{n})$. To obtain our speed-up, we relax the conditions on the augmenting paths and compute, in each phase, a set of augmenting paths that are not restricted to be shortest or vertex-disjoint. As a result, our algorithm computes substantially more augmenting paths in each phase, reducing the number of phases from $O(\sqrt{n})$ to $O(n^{2/5})$. By using small vertex separators, the execution of each phase takes $\tilde{O}(m)$ time on average. For planar graphs, we combine our algorithm with efficient shortest path data structures to obtain a minimum-cost perfect matching in $\tilde{O}(n^{6/5} \log (nC))$ time. This improves upon the recent $\tilde{O}(n^{4/3} \log (nC))$ time algorithm by Asathulla et al. [SODA 2018].

Nathaniel Lahn, Sharath Raghvendra
Virginia Tech
lahnn@vt.edu, sharathr@vt.edu

CP10

Flow-Cut Gaps and Face Covers in Planar Graphs

The relationship between the sparsest cut and the maximum concurrent multi-flow in graphs has been studied extensively. For general graphs with k terminal pairs, the flow-cut gap is $O(\log k)$, and this is tight. For planar graphs, it has been conjectured that their flow-cut gap is $O(1)$, while the known bounds place the gap somewhere between 2 (Lee and Raghavendra, 2003) and $O(\sqrt{\log k})$ (Rao, 1999). Okamura and Seymour (1981) show that when all the terminals of a planar network lie on a single face, the flow-cut gap is exactly 1. This setting can be generalized by considering planar networks where the terminals lie on $\gamma > 1$ faces in some fixed planar drawing. Lee and Sidiropoulos (2009) proved that the flow-cut gap is bounded by a function of γ, and Chekuri, Shepherd, and Weibel (2013) showed that the gap is at most 3γ. We prove that the flow-cut gap is $O(\log \gamma)$, by showing that the edge-weighted shortest-path metric induced on the terminals admits a stochastic embedding into trees with distortion $O(\log \gamma)$, which is tight. For *vertex-capacitated networks,* Lee, Mendel, and Moharrami (2015) showed that the vertex-capacitated flow-cut gap is $O(1)$ on planar networks whose terminals lie on a single face. We prove that the flow-cut gap is $O(\gamma)$ for vertex-capacitated instances when the terminals lie on $\gamma > 1$ faces. In fact, this result holds in the more general setting of submodular vertex capacities.

Havana Rika, Robert Krauthgamer
Weizmann Institute of Science
havana.rika@weizmann.ac.il,
robert.krauthgamer@weizmann.ac.il

James R. Lee
University of Washington
jrl@cs.washington.edu

CP10

On Constant Multi-Commodity Flow-Cut Gaps for Families of Directed Minor-Free Graphs

The multi-commodity flow-cut gap is a fundamental parameter that affects the performance of several divide & conquer algorithms, and has been extensively studied for various classes of undirected graphs. It has been shown by Linial, London and Rabinovich [1994] and by Aumann and Rabani [1998] that for general n-vertex graphs it is bounded by $O(\log n)$ and the Gupta-Newman-Rabinovich-Sinclair conjecture [2004] asserts that it is $O(1)$ for any family of graphs that excludes some fixed minor. The flow-cut gap is poorly understood for the case of directed graphs. We show that for uniform demands it is $O(1)$ on directed series-parallel graphs, and on directed graphs of bounded pathwidth. These are the first constant upper bounds of this type for some non-trivial family of directed graphs. We also obtain $O(1)$ upper bounds for the general multi-commodity flow-cut gap on directed trees and cycles. These bounds are obtained via new embeddings and Lipschitz quasipartitions for quasimetric spaces, which generalize analogous results form the metric case, and could be of independent interest. Finally, we discuss limitations of methods that were developed for undirected graphs, such as random partitions, and random embeddings.

Ario Salmasi
The Ohio State University
ario.salmasi@gmail.com

Anastasios Sidiropoulos
University of Illinois at Chicago
sidiropo@uic.edu

Vijay Sridhar
Mathworks
vijaysridhar1991@gmail.com

CP10

Maximum Integer Flows in Directed Planar Graphs with Vertex Capacities and Multiple Sources and Sinks

We consider the problem of finding maximum flows in planar graphs with capacities on both vertices and arcs and with multiple sources and sinks. We present three algorithms when the capacities are integers. The first algorithm runs in $O(n \log^3 n + kn)$ time when all capacities are bounded, where n is the number of vertices in the graph and k is the number of terminals. This algorithm is the first to solve the vertex-disjoint paths problem in near-linear time when k is bounded but larger than 2. The second algorithm runs in $O(k^5 n \operatorname{polylog}(nU))$ time, where U is the largest finite capacity of a single vertex. Finally, when $k = 3$, we present an algorithm that runs in $O(n \log n)$ time; this algorithm works even when the capacities are arbitrary reals. Our algorithms improve on the fastest previously known algorithms when k is bounded and U is bounded by a polynomial in n. Prior to this result, the fastest algorithms ran in $O(n^2/\log n)$ time for real capacities, $O(n^{3/2} \log n \log U)$ for integer capacities, and $\tilde{O}(n^{10/7})$ for unit capacities, even when $k = 3$.

Yipu Wang
University of Illinois at Urbana-Champaign

ywang298@illinois.edu

CP11

Quantum Algorithms and Approximating Polynomials for Composed Functions with Shared Inputs

We give faster quantum algorithms for evaluating composed functions whose inputs may be shared between bottom-level gates. Let f be a Boolean function and consider a function F obtained by applying f to conjunctions of possibly overlapping subsets of n variables. If f has quantum query complexity $Q(f)$, we give an algorithm for evaluating F using $\tilde{O}(\sqrt{Q(f)\cdot n})$ quantum queries. This improves on the bound of $O(Q(f)\cdot\sqrt{n})$ that follows by treating each conjunction independently, and is tight for worst-case choices of f. Using completely different techniques, we prove a similar tight composition theorem for the approximate degree of f. By recursively applying our composition theorems, we obtain a nearly optimal $\tilde{O}(n^{1-2^{-d}})$ upper bound on the quantum query complexity and approximate degree of linear-size depth-d AC^0 circuits. As a consequence, such circuits can be PAC learned in subexponential time, even in the challenging agnostic setting. Prior to our work, a subexponential-time algorithm was not known even for linear-size depth-3 AC^0 circuits. We also show that any substantially faster learning algorithm will require fundamentally new techniques.

Mark Bun
Simons Institute for the Theory of Computing
markmbun@gmail.com

Robin Kothari
Microsoft Research
robin@robinkothari.com

Justin Thaler
Georgetown University
justin.thaler@georgetown.edu

CP11

Reproducibility and Pseudo-determinism in Log-space

A curious property of randomized log-space search algorithms is that their outputs are often longer than their workspace. This leads to the question: how can we reproduce the results of a randomized log space computation without storing the output or randomness verbatim? Running the algorithm again with new random bits may result in a new (and potentially different) output. We show that every problem in search-RL has a randomized log-space algorithm where the output can be reproduced. Specifically, we show that for every problem in search-RL, there are a pair of log-space randomized algorithms A and B where for every input x, A will output some string t_x of size $O(\log n)$, such that B when running on (x, t_x) will be pseudo-deterministic: that is, running B multiple times on the same input (x, t_x) will result in the same output on all executions with high probability. Thus, by storing only $O(\log n)$ bits in memory, it is possible to reproduce the output of a randomized log-space algorithm. An algorithm is reproducible without storing *any* bits in memory (i.e., $|t_x| = 0$) if and only if it is pseudo-deterministic. We show pseudo-deterministic algorithms for finding paths in undirected graphs and Eulerian graphs using logarithmic space. Our algorithms are substantially faster than the best known deterministic algorithms for finding paths in such graphs in log-space.

Yang Liu
Stanford University
yangpatil@gmail.com

Ofer Grossman
Berkeley
ofer.grossman@gmail.com

CP11

A Deterministic PTAS for the Algebraic Rank of Bounded Degree Polynomials

We consider the problem of computing the algebraic rank of a given set of multivariate polynomials over a field of characteristic zero. The notion of algebraic rank naturally generalizes the notion of rank in linear algebra. The decision version of the problem is equivalent to the celebrated *Polynomial Identity Testing* (PIT) problem. A randomized polynomial time algorithm for this problem is known thanks to the classical Jacobian criterion, and the Schwartz-Zippel lemma. However, no deterministic algorithm is known. In fact, a deterministic algorithm would imply a deterministic algorithm for the PIT problem which itself would imply circuit complexity lower bounds (Kabanets, Impagliazzo, STOC'03). We present a deterministic polynomial time approximation scheme (PTAS) for computing the algebraic rank of a set of bounded degree polynomials. More specifically, we give an algorithm that takes as input a set $\mathbf{f} := \{f_1, \ldots, f_n\} \subset \mathbb{F}[x_1, \ldots, x_m]$ of polynomials with degrees bounded by d, and a rational number $\epsilon > 0$ and runs for time $O((\frac{nmd}{\epsilon})^{O(\frac{d^2}{\epsilon})} \cdot M(n))$, where $M(n)$ is the time required to compute the rank of an $n \times n$ matrix (with field entries), and finally outputs a number r, such that r is at least $(1-\epsilon)$ times the algebraic rank of $\mathbf{f}$.

Vishwas Bhargava
Rutgers University
vishwas1384@gmail.com

Markus Bläser
Saarland University
mblaeser@cs.uni-saarland.de

Gorav Jindal
Aalto University
gorav.jindal@gmail.com

Anurag Pandey
Max-Planck-Institut fuer Informatik
apandey@mpi-inf.mpg.de

CP11

Pseudorandomness for Read-k DNF Formulas

The design of pseudorandom generators and deterministic approximate counting algorithms for DNF formulas are important challenges in unconditional derandomization. Numerous works on these problems have focused on the subclass of *small-read* DNF formulas, which are formulas in which each variable occurs a bounded number of times. Our first main result is a pseudorandom generator which ε-fools M-term read-k DNFs using seed length $\mathrm{poly}(k, \log(1/\varepsilon)) \cdot \log M + O(\log n)$. This seed length is exponentially shorter, as a function of both k and $1/\varepsilon$, than the best previous PRG for read-k DNFs. We also give a

deterministic algorithm that approximates the number of satisfying assignments of an M-term read-k DNF to any desired $(1+\varepsilon)$-multiplicative accuracy in time

$$\mathrm{poly}(n)\cdot\min\left\{(M/\varepsilon)^{\mathrm{poly}(k,\log(k/\varepsilon))},(M/\varepsilon)^{\tilde{O}(\log((k\log M)/\varepsilon))}\right\}.$$

The common essential ingredients in these pseudorandomness results are new analytic inequalities for read-k DNFs. These inequalities may be of independent interest and utility; as an example application, we use them to obtain a significant improvement on the previous state of the art for agnostically learning read-k DNFs.

Rocco A. Servedio, Li-Yang Tan
Columbia University
rocco@cs.columbia.edu, liyang@cs.stanford.edu

CP11

Near-optimal Bootstrapping of Hitting Sets for Algebraic Circuits

The classical lemma of Ore-DeMillo-Lipton-Schwartz-Zippel states that any nonzero n-variate degree-s polynomial f, will evaluate to a nonzero value at some point on a grid S^n with $|S|>s$. Thus, there is a deterministic polynomial identity test (PIT) for all degree-s size-s algebraic circuits in n variables that runs in time $poly(s)(s+1)^n$. In a surprising recent result, Agrawal, Ghosh and Saxena (STOC'18) showed any deterministic blackbox PIT algorithm for degree-s, size-s, n-variate circuits with running time as bad as $s^{n^{0.5-\delta}}$, where $\delta>0$, can be used to construct blackbox PIT algorithms for degree-s size s circuits with running time $s^{\exp(\exp(O(\log^* s)))}$. The authors asked if a similar conclusion followed if their hypothesis was weakened to having deterministic PIT with running time $s^{o(n)}$. In this paper, we answer their question in the affirmative. We show that, given a deterministic blackbox PIT that runs in time $(s^{o(n)})H(n)$, where $H(n)$ is an arbitrary function, for all degree-s size-s algebraic circuits over n variables, we can obtain a deterministic blackbox PIT that runs in time $s^{\exp(\exp(O(\log^* s)))}$ for all degree-s size-s algebraic circuits over n variables. In other words, any blackbox PIT with just a slightly non-trivial exponent of s compared to the trivial $s^{O(n)}$ test can be used to give a nearly polynomial time blackbox PIT algorithm.

Mrinal Kumar
Simons Institute for the Theory of Computing, Berkeley, USA
mrinalkumar08@gmail.com

Ramprasad Saptharishi
Tata Institute of Fundamental Research, Mumbai, India
ramprasad@tifr.res.in

Anamay Tengse
Tata Institute of Fundamental Research, Mumbai
tengse.anamay@tifr.res.in

CP12

A Ptas for ℓ_p-Low Rank Approximation

A number of recent works have studied algorithms for entrywise ℓ_p-low rank approximation, namely, algorithms which given an $n\times d$ matrix A (with $n\geq d$), output a rank-k matrix B minimizing $\|A-B\|_p^p=\sum_{i,j}|A_{i,j}-B_{i,j}|^p$ when $p>0$; and $\|A-B\|_0=\sum_{i,j}[A_{i,j}\neq B_{i,j}]$ for $p=0$. On the algorithmic side, for $p\in(0,2)$, we give the first $(1+\epsilon)$-approximation algorithm running in time $n^{\mathrm{poly}(k/\epsilon)}$. Further, for $p=0$, we give the first almost-linear time approximation scheme for what we call the Generalized Binary ℓ_0-Rank-k problem. Our algorithm computes $(1+\epsilon)$-approximation in time $(1/\epsilon)^{2^{O(k)}/\epsilon^2}\cdot nd^{1+o(1)}$. On the hardness of approximation side, for $p\in(1,2)$, assuming the Small Set Expansion Hypothesis and the Exponential Time Hypothesis (ETH), we show that there exists $\delta:=\delta(\alpha)>0$ such that the entrywise ℓ_p-Rank-k problem has no α-approximation algorithm running in time 2^{k^δ}.

Frank Ban
UC Berkeley
fban@berkeley.edu

Vijay Bhattiprolu
CMU
vpb@cs.cmu.edu

Karl Bringmann
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
kbringma@mpi-inf.mpg.de

Pavel Kolev
MPI-INF
pavelkolevbg@gmail.com

Euiwoong Lee
Carnegie Mellon University
euiwoong@cims.nyu.edu

David Woodruff
CMU
dwoodruf@andrew.cmu.edu

CP12

Every Testable (Infinite) Property of Bounded-degree Graphs Contains An Infinite Hyperfinite Subproperty

One of the most fundamental questions in graph property testing is to characterize the combinatorial structure of properties that are testable with a constant number of queries. We work towards an answer to this question for the bounded-degree graph model, where the input graphs have maximum degree bounded by a constant d. In this model, it is known (among other results) that every *hyperfinite* property is constant-query testable, where, informally, a graph property is hyperfinite, if for every $\delta>0$ every graph in the property can be partitioned into small connected components by removing δn edges. In this paper we show that hyperfiniteness plays a role in *every* testable property, i.e. we show that every testable property is either finite (which trivially implies hyperfiniteness and testability) or contains an infinite hyperfinite subproperty. A simple consequence of our result is that no infinite graph property that only consists of expander graphs is constant-query testable. Based on the above findings, one could ask if every infinite testable non-hyperfinite property might contain an infinite family of expander (or near-expander) graphs. We show that this is not true.

Hendrik Fichtenberger
TU Dortmund
hendrik.fichtenberger@tu-dortmund.de

Pan Peng

University of Sheffield
p.peng@sheffield.ac.uk

Christian Sohler
Department of Computer Science
Technische Universitat Dortmund
christian.sohler@tu-dortmund.de

CP12

Testing Halfspaces over Rotation-invariant Distributions

We present an algorithm for testing halfspaces over arbitrary, unknown rotation-invariant distributions. Using $\widetilde{O}(\sqrt{n}\epsilon^{-7})$ random examples of an unknown function f, the algorithm determines with high probability whether f is of the form $f(x) = sign(\sum_i w_i x_i - t)$ or is ϵ-far from all such functions. This sample size is significantly smaller than the well-known requirement of $\Omega(n)$ samples for learning halfspaces, and known lower bounds imply that our sample size is optimal up to logarithmic factors. The algorithm is distribution-free in the sense that it requires no knowledge of the distribution aside from the promise of rotation invariance. To prove the correctness of this algorithm we present a theorem relating the distance between a function and a halfspace to the distance between their centers of mass, that applies to arbitrary distributions.

Nathaniel Harms
University of Waterloo
nharms@uwaterloo.ca

CP12

Anaconda: A Non-adaptive Conditional Sampling Algorithm for Distribution Testing

We investigate distribution testing with access to non-adaptive conditional samples. In the conditional sampling model, the algorithm is given the following access to a distribution: it submits a query set S to an oracle, which returns a sample from the distribution conditioned on being from S. In the non-adaptive setting, all query sets must be specified in advance of viewing the outcomes. Our main result is the first polylogarithmic-query algorithm for equivalence testing, deciding whether two unknown distributions are equal to or far from each other. This is an exponential improvement over the previous best upper bound, and demonstrates that the complexity of the problem in this model is intermediate to the the complexity of the problem in the standard sampling model and the adaptive conditional sampling model. We also significantly improve the sample complexity for the easier problems of uniformity and identity testing. For the former, our algorithm requires only $\tilde{O}(\log n)$ queries, matching the information-theoretic lower bound up to a $O(\log\log n)$-factor. Our algorithm works by reducing the problem from ℓ_1-testing to ℓ_∞-testing, which enjoys a much cheaper sample complexity. Necessitated by the limited power of the non-adaptive model, our algorithm is very simple to state. However, there are significant challenges in the analysis, due to the complex structure of how two arbitrary distributions may differ.

Gautam Kamath
MIT
g@csail.mit.edu

Christos Tzamos
University of Wisconsin, Madison
tzamos@wisc.edu

CP12

Testing Matrix Rank, Optimally

We show that for the problem of testing if a matrix $A \in F^{n\times n}$ has rank at most d, or requires changing an ϵ-fraction of entries to have rank at most d, there is a *non-adaptive* query algorithm making $\widetilde{O}(d^2/\epsilon)$ queries. Our algorithm works for any field F. This improves upon the previous $O(d^2/\epsilon^2)$ bound (SODA'03), and bypasses an $\Omega(d^2/\epsilon^2)$ lower bound of (KDD'14) which holds if the algorithm is required to read a submatrix. Our algorithm is the first such algorithm which does not read a submatrix, and instead reads a carefully selected non-adaptive pattern of entries in rows and columns of A. We complement our algorithm with a matching query complexity lower bound for non-adaptive testers over any field. We also give tight bounds of $\widetilde{\Theta}(d^2)$ queries in the sensing model for which query access comes in the form of $\langle X_i, A\rangle := tr(X_i^\top A)$; perhaps surprisingly these bounds do not depend on ϵ. We next develop a novel property testing framework for testing numerical properties of a real-valued matrix A more generally, which includes the stable rank, Schatten-p norms, and SVD entropy. Specifically, we propose a bounded entry model, where A is required to have entries bounded by 1 in absolute value. We give upper and lower bounds for a wide range of problems in this model, and discuss connections to the sensing model above.

Maria-Florina Balcan
Carnegie Mellon University
ninamf@cs.cmu.edu

Yi Li
Nanyang Technological University
yili@ntu.edu.sg

David Woodruff, Hongyang Zhang
Carnegie Mellon University
dwoodruf@cs.cmu.edu, hongyanz@cs.cmu.edu

CP13

Approximation of Trees by Self-nested Trees

The class of self-nested trees presents remarkable compression properties because of the systematic repetition of subtrees in their structure. In this paper, we provide a better combinatorial characterization of this specific family of trees. In particular, we show from both theoretical and practical viewpoints that complex queries can be quickly answered in self-nested trees compared to general trees. We also present an approximation algorithm of a tree by a self-nested one that can be used in fast prediction of edit distance between two trees.

Romain Azais
Inria
romain.azais@inria.fr

Jean-Baptiste Durand
Université Grenoble Alpes
jean-baptiste.durand@univ-grenoble-alpes.fr

Christophe Godin
Inria

christophe.godin@inria.fr

CP13

Simple and Fast BlockQuicksort using Lomuto's Partitioning Scheme

This paper presents simple variants of the BlockQuicksort algorithm described by Edelkamp and Weiss (ESA 2016). The simplification is achieved by using Lomuto's partitioning scheme instead of Hoare's crossing pointer technique to partition the input. To achieve a robust sorting algorithm that works well on many different input types, the paper introduces a novel two-pivot variant of Lomuto's partitioning scheme. A surprisingly simple twist to the generic two-pivot quicksort approach makes the algorithm robust. The paper provides an analysis of the theoretical properties of the proposed algorithms and compares them to their competitors. The analysis shows that Lomuto-based approaches incur a higher average sorting cost than the Hoare-based approach of BlockQuicksort. Moreover, the analysis is particularly useful to reason about pivot choices that suit the two-pivot approach. An extensive experimental study shows that, despite their worse theoretical behavior, the simpler variants perform as well as the original version of BlockQuicksort.

Martin Aumüller
IT University of Copenhagen
maau@itu.dk

Nikolaj Hass
IT University Of Copenhagen
niah@itu.dk

CP13

Lightweight Distributed Suffix Array Construction

We present two new distributed suffix array construction algorithms. One of our algorithms requires only half as much memory as its competitor (PSAC) [Flick & Aluru, SC 2015], while achieving similar speed. In practice, we can compute on the same hardware suffix arrays for text twice as large as PSAC. The other algorithm still requires less memory than PSAC but is faster on some instances. As a by-product, we also engineered the first distributed string sorting algorithm. All of our algorithms are tested on text collections of up to 115 GB and running on 1280 cores.

Florian Kurpicz, Johannes Fischer
TU Dortmund
florian.kurpicz@tu-dortmund.de, johannes.fischer@cs.tu-dortmund.de

CP13

Worst-case Efficient Sorting with QuickMergesort

The two most prominent solutions for the sorting problem are Quicksort and Mergesort. While Quicksort is very fast on average, Mergesort additionally gives worst-case guarantees, but needs extra space for a linear number of elements. Worst-case efficient in-place sorting, however, remains a challenge: the standard solution, Heapsort, suffers from a bad cache behavior and is also not overly fast for in-cache instances. In this work we present median-of-medians QuickMergesort (MoMQuickMergesort), a new variant of QuickMergesort, which combines Quicksort with Mergesort allowing the latter to be implemented in place. Our new variant applies the median-of-medians algorithm for selecting pivots in order to circumvent the quadratic worst case. Indeed, we show that it uses at most $n \log n + 1.6n$ comparisons for n large enough. We experimentally confirm the theoretical estimates and show that the new algorithm outperforms Heapsort by far and is only around 10% slower than Introsort (std::sort implementation of stdlibc++), which has a rather poor guarantee for the worst case. We also simulate the worst case, which is only around 10% slower than the average case. In particular, the new algorithm is a natural candidate to replace Heapsort as a worst-case stopper in Introsort.

Armin Weiß
Universität Stuttgart
armin.weiss@fmi.uni-stuttgart.de

Stefan Edelkamp
King's College London
stefan.edelkamp@gmail.com

CP14

Sublinear Algorithms for (Delta + 1) Vertex Coloring

Any graph with maximum degree Δ admits a proper vertex coloring with $\Delta + 1$ colors that can be found via a simple sequential greedy algorithm in linear time and space. But can one find such a coloring via a sublinear algorithm? We answer this fundamental question in the affirmative for several canonical classes of sublinear algorithms including graph streaming, sublinear time, and massively parallel computation (MPC) algorithms. In particular, we design (a) a single-pass semi-streaming algorithm in dynamic streams using $\tilde{O}(n)$ space, (b) a sublinear-time algorithm in the standard query model that allows neighbor queries and pair queries using $\tilde{O}(n\sqrt{n})$ time (which we also prove is optimal), and (c) a parallel algorithm in the massively parallel computation (MPC) model using $\tilde{O}(n)$ memory per machine and $O(1)$ MPC rounds. At the core of our results is a remarkably simple meta-algorithm for the $(\Delta+1)$ coloring problem: Sample $O(\log n)$ colors for each vertex uniformly at random from the $\Delta + 1$ colors; find a proper coloring of the graph using the sampled colors. As our main result, we prove that the sampled set of colors with high probability contains a proper coloring of the input graph. The sublinear algorithms are then obtained by designing efficient algorithms for finding a proper coloring of the graph from the sampled colors in the corresponding models.

Sepehr Assadi, Yu Chen, Sanjeev Khanna
University of Pennsylvania
sassadi@cis.upenn.edu, chenyu2@cis.upenn.edu, sanjeev@cis.upenn.edu

CP14

Optimal Distributed Coloring Algorithms for Planar Graphs in the Local Model

In this paper, we consider distributed coloring for planar graphs with a small number of colors. Our main result is an optimal (up to a constant factor) $O(\log n)$ time algorithm for 6-coloring planar graphs. Our algorithm is based on a novel technique that in a nutshell detects small structures that can be easily colored given a proper coloring of the rest of the vertices and removes them from the graph until the graph contains a small enough number of edges.

We believe this technique might be of independent interest. In addition, we present a lower bound for 4-coloring planar graphs that essentially shows that any algorithm (deterministic or randomized) for 4-coloring planar graphs requires $\Omega(n)$ rounds. We therefore completely resolve the problems of 4-coloring and 6-coloring for planar graphs in the LOCAL model.

Shiri Chechik
Tel-Aviv University, Israel
shiri.chechik@gmail.com

Doron Mukhtar
Tel-Aviv University
doron.muk@gmail.com

CP14

Distributed Maximal Independent Set Using Small Messagesu

Abstract not available

Mohsen Ghaffari
ETH Zurich
ghaffari@inf.ethz.ch

CP14

Oblivious Resampling Oracles and Parallel Algorithms for the Lopsided Lovász Local Lemma

The Lovász Local Lemma (LLL) is a probabilistic tool which shows that, if a collection of "bad' events $\mathcal{B}$ in a probability space are not too likely and not too interdependent, then there is a positive probability that no bad-events in $\mathcal{B}$ occur. Moser & Tardos (2010) gave sequential and parallel algorithms which transformed most applications of the variable-assignment LLL into efficient algorithms. A framework of Harvey & Vondrák (2015) based on "resampling oracles' extended this give very general sequential algorithms for other probability spaces satisfying the Lopsided Lovász Local Lemma (LLLL). We describe a new structural property of resampling oracles which holds for all known resampling oracles, which we call "obliviousness.' Essentially, it means that the interaction between two bad-events B, B' depends only on the randomness used to resample B. This has two major consequences. First, it is the key to achieving a unified parallel LLLL algorithm, giving the first RNC algorithms for rainbow perfect matchings and rainbow hamiltonian cycles of K_n. Second, this property allows us to build LLLL probability spaces out of a relatively simple "atomic' set of events. Using this framework, we get the first sequential resampling oracle for rainbow perfect matchings on the complete s-uniform hypergraph $K_n^{(s)}$, and the first commutative resampling oracle for hamiltonian cycles of K_n.

David G. Harris
University of Maryland
davidgharris29@gmail.com

CP14

Distributed Triangle Detection via Expander Decomposition

We present improved distributed algorithms for triangle detection and its variants in the CONGEST model. We show that Triangle Detection, Counting, and Enumeration can be solved in $\tilde{O}(n^{1/2})$ rounds. In contrast, the previous state-of-the-art bounds for Triangle Detection and Enumeration were $\tilde{O}(n^{2/3})$ and $\tilde{O}(n^{3/4})$, respectively, due to Izumi and Le Gall (PODC 2017). The main novelty in this work is a distributed graph partitioning algorithm. We show that in $\tilde{O}(n^{1-\delta})$ rounds we can partition the edge set of the network $G = (V, E)$ into three parts $E = E_m \cup E_s \cup E_r$ such that

- Each connected component induced by E_m has minimum degree $\Omega(n^\delta)$ and conductance $\Omega(1/\text{poly}\log(n))$.
- The subgraph induced by E_s has arboricity at most n^δ.
- $|E_r| \leq |E|/6$.

All our algorithms are based on the following generic framework, which we believe is of interest beyond this work. Roughly, we deal with E_s by an algorithm that is efficient for low-arboricity graphs, and deal with E_r using recursive calls. For each connected component induced by E_m, we are able to simulate congested clique algorithms with small overhead by applying a routing algorithm due to Ghaffari, Kuhn, and Su (PODC 2017) for high conductance graphs.

Yi-Jun Chang, Seth Pettie
University of Michigan
cyijun@umich.edu, pettie@umich.edu

Hengjie Zhang
Tsinghua University
zhanghj15@mails.tsinghua.edu.cn

CP15

Improving the Smoothed Complexity of Flip for Max-cut Problems

Finding locally optimal solutions for max-cut and max-k-cut are well-known PLS-complete problems. An instinctive approach to finding such solutions is the FLIP method. Even though FLIP requires exponential time in worst-case instances, it tends to terminate quickly in practical instances. To explain this discrepancy, the run-time of FLIP has been studied in the smoothed complexity framework. Etscheid and Röglin (ACM Trans. Algorithms, 2017) showed that the smoothed complexity of FLIP for max-cut is quasi-polynomial. Angel, Bubeck, Peres and Wei (STOC 2017) showed that the smoothed complexity of FLIP for max-cut in complete graphs is $O(\phi^5 n^{15.1})$, where ϕ is an upper bound on the random edge-weight density and n is the number of vertices in the input graph. In this work, we improve the run-time bound for complete graphs. We prove that the smoothed complexity of FLIP in complete graphs is $O(\phi n^{7.83})$. Our results are based on a carefully chosen matrix whose rank captures the run-time of the method along with improved rank bounds for this matrix and an improved union bound based on this matrix. In addition, our techniques provide a general framework for analyzing FLIP in the smoothed framework. We illustrate this general framework by showing that the smoothed complexity of FLIP for max-3-cut in complete graphs is polynomial and for max-k-cut in arbitrary graphs is quasi-polynomial.

Ali Bibak
University of Illinois, Urbana-Champaign
bibakse2@illinois.edu

Charles A. Carlson
University of Colorado Boulder
chca0914@colorado.edu

Karthekeyan Chandrasekaran
Department of Industrial and Enterprise Systems Engineering
University of Illinois, Urbana-Champaign
karthe@illinois.edu

CP15

On the Number of Circuits in Regular Matroids (with Connections to Lattices and Codes)

We show that for any regular matroid on m elements and any $\alpha \geq 1$, the number of α-minimum circuits, or circuits whose size is at most an α-multiple of the minimum size of a circuit in the matroid is bounded by $m^{O(\alpha^2)}$. This generalizes a result of Karger for the number of α-minimum cuts in a graph. As a consequence, we obtain similar bounds on the number of α-shortest vectors in "totally unimodular" lattices and on the number of α-minimum weight codewords in "regular" codes.

Rohit Gurjar
Indian Institute of Technology Bombay
rohitgurjar0@gmail.com

Nisheeth K. Vishnoi
EPFL
EPFL
nisheeth.vishnoi@gmail.com

CP15

Nash Flows over Time with Spillback

Modeling traffic in road networks is a widely studied but challenging problem, especially under the assumption that drivers act selfishly. A common approach used in simulation software is the deterministic queuing model, for which the structure of dynamic equilibria has been studied extensively in the last couple of years. The basic idea is to model traffic by a continuous flow that travels over time from a source to a sink through a network, in which the arcs are endowed with transit times and capacities. Whenever the flow rate exceeds the capacity a queue builds up and the infinitesimally small flow particles wait in line in front of the bottleneck. It was not possible, until now, to represent spillback in this model, which was a big drawback, since spillback has a huge impact on travel times in highly congested regions. We extend the model by introducing a storage capacity that bounds the total amount of flow on each arc. If an arc gets full, the inflow capacity is reduced to the current outflow rate, which can cause queues on previous arcs, i.e., spillback. We carry over the main results of the original model to our generalization and characterize dynamic equilibria, called Nash flows over time, by sequences of particular static flows, we call spillback thin flows. Furthermore, we give a constructive proof for the existence of dynamic equilibria, which suggests an algorithm for their computation. This solves an open problem stated by Koch and Skutella in 2010.

Leon Sering
Institute of Mathematics - Technische Universität Berlin
sering@math.tu-berlin.de

Laura Vargas Koch
School of Business and Economics - RWTH Aachen University
laura.vargas@oms.rwth-aachen.de

CP15

Computing all Wardrop Equilibria Parametrized by the Flow Demand

We develop an algorithm that computes for a given undirected or directed network with flow-dependent piece-wise linear edge cost functions *all* Wardrop equilibria as a function of the flow demand. Our algorithm is based on Katzenelson's homotopy method for electrical networks. The algorithm uses a bijection between vertex potentials and flow excess vectors that is piecewise linear in the potential space and where each linear segment can be interpreted as an augmenting flow in a residual network. The algorithm iteratively increases the excess of one or more vertex pairs until the bijection reaches a point of non-differentiability. Then, the next linear region is chosen in a simplex-like pivot step and the algorithm proceeds. We first show that this algorithm correctly computes all Wardrop equilibria in undirected single-commodity networks along the chosen path of excess vectors. We then adapt our algorithm to also work for discontinuous cost functions which allows to model directed edges and/or edge capacities. Our algorithm is output-polynomial in non-degenerate instances where the solution curve never hits a point where the cost function of more than one edge becomes non-differentiable. For degenerate instances we still obtain an output-polynomial algorithm computing the linear segments of the bijection by a convex program. The latter technique also allows to handle multiple commodities.

Philipp Warode, Max Klimm
School of Business and Economics, HU Berlin
philipp.warode@hu-berlin.de, max.klimm@hu-berlin.de

CP15

Minimum Cut and Minimum k-Cut in Hypergraphs via Branching Contractions

Random edge contractions (Karger, SODA '93) provide a particularly simple and elegant method for computing the minimum cut of a graph. Recent work (Chandrasekharan, Xu, and Yu, SODA '18; Ghaffari, Karger, and Panigrahi, SODA '17) has extended the random contraction technique from graphs to finding minimum cuts and minimum k-cuts in hypergraphs. We further this line of work by providing a *branching* randomized contraction technique for hypergraphs. Our algorithms perform branches as a random response to the size of a hyperedge selected for contraction. The analysis then hinges on being able to quantify the (random) structure of the resulting computation tree and the randomness of the contracted hypergraph. This new technique leads to the following improvements in running time (m is the number of hyperedges, n the number of vertices, and p the total size of all hyperedges):

- We give an algorithm that runs in $\widetilde{O}\left(mn^{2k-2}\right)$ time for finding a minimum k-cut in hypergraphs of arbitrary rank (maximum size of a hyperedge). This improves the best known running time for $k > 2$.
- We give another algorithm that runs in $\widetilde{O}\left(n^{\max\{r,2k-2\}}\right)$ time for finding a minimum k-cut in hypergraphs of constant rank r. This betters the best known running times for dense hypergraphs.

Our techniques and results extend to the problems of minimum hedge-cut and minimum hedge-k-cut as well.

Kyle Fox

Department of Computer Science
The University of Texas at Dallas
kyle.fox@utdallas.edu

Debmalya Panigrahi
Duke University
debmalya@cs.duke.edu

Fred Zhang
Harvard University
hzhang@g.harvard.edu

CP16

On the Rank of a Random Sparse Binary Matrix

We study the rank of a random $n \times m$ matrix $\mathbf{A}_{n,m;k}$ with entries from $GF(2)$, and exactly k unit entries in each column, the other entries being zero. The columns are chosen independently and uniformly at random from the set of all $\binom{n}{k}$ such columns. We obtain an asymptotically correct estimate for the rank as a function of the number of columns m in terms of c, n, k, and where $m = cn/k$. The matrix $\mathbf{A}_{n,m;k}$ forms the vertex-edge incidence matrix of a k-uniform random hypergraph H. The rank of $\mathbf{A}_{n,m;k}$ can be expressed as follows. Let $|C_2|$ be the number of vertices of the 2-core of H, and $|E(C_2)|$ the number of edges. Let m^* be the value of m for which $|C_2| = |E(C_2)|$. Then w.h.p. for $m < m^*$ the rank of $A_{n,m;k}$ is asymptotic to m, and for $m \geq m^*$ the rank is asymptotic to $m - |E(C_2)| + |C_2|$. In addition, assign i.i.d. $U[0,1]$ weights $X_i, i \in \{1, 2, ...m\}$ to the columns, and define the weight of a set of columns S as $X(S) = \sum_{j \in S} X_j$. Define a basis as a set of $n-1_k$ even linearly independent columns. We obtain an asymptotically correct estimate for the minimum weight basis. This generalises the well-known result of Frieze [On the value of a random minimum spanning tree problem, Discrete Applied Mathematics, (1985)] that, for $k = 2$, the expected length of a minimum weight spanning tree tends to $\zeta(3) \sim 1.202$.

Alan Frieze
Carnegie Mellon University
Department of Mathematical Sciences
alan@random.math.cmu.edu

Colin Cooper
Kings College, London
colin.cooper@kcl.ac.uk

Wesley Pegden
Carnegie Mellon University
wes@math.cmu.edu

CP16

Rapid Mixing of the Switch Markov Chain for Strongly Stable Degree Sequences and 2-Class Joint Degree Matrices

The *switch Markov chain* has been extensively studied as the most natural Markov Chain Monte Carlo approach for sampling graphs with prescribed degree sequences. We use comparison arguments to show that the switch chain mixes rapidly in two different settings. We first study the classic problem of uniformly sampling simple undirected, as well as bipartite, graphs with a given degree sequence. We apply an embedding argument, involving a Markov chain defined by Jerrum and Sinclair (TCS, 1990) for sampling graphs that *almost* have a given degree sequence, to show rapid mixing for degree sequences satisfying *strong stability*, a notion closely related to P-stability. This results in a much shorter proof that unifies and extends the currently known rapid mixing results for the switch chain. In particular, our work resolves an open problem posed by Greenhill (SODA, 2015). Secondly, in order to illustrate the power of our approach, we study the problem of uniformly sampling graphs for which—in addition to the degree sequence—a *joint degree distribution* is given. Although the problem was formalized over a decade ago, small progress has been made on the random sampling of such graphs. The case of a single degree class reduces to sampling of regular graphs, but beyond this almost nothing is known. We fully resolve the case of two degree classes, by showing that the switch Markov chain is *always* rapidly mixing.

Georgios Amanatidis, Pieter Kleer
Centrum Wiskunde & Informatica (CWI)
amana284@gmail.com, kleer@cwi.nl

CP16

On Coalescence Time in Graphs-When Is Coalescing as Fast as Meeting?

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. We provide a powerful toolkit that results in tight bounds for various topologies. As a general result, we establish that for graphs whose meeting time is only marginally larger than the mixing time, the coalescence time of n random walks equals the meeting time up to constant factors. We show that this is tight. 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 $O(n^3)$ (which is tight for the Barbell graph); surprisingly even such a basic question about the coalescing time was not answered before this work. By duality, our results give bounds on the voter model and therefore give bounds on the consensus time in arbitrary undirected graphs. We also establish a new bound on the hitting time and cover time of regular graphs, improving and tightening previous results by Broder and Karlin, as well as those by Aldous and Fill.

Varun Kanade
University of Oxford
varunk@cs.ox.ac.uk

Frederik Mallmann-Trenn
MIT
fmallman@sfu.ca

Thomas Sauerwald
University of Cambridge
thomas.sauerwald@cl.cam.ac.uk

CP16

Xor Codes and Learning Sparse Parities with Noise

A k-LIN instance is a system of m equations over n variables of the form $s_{i_1} + \cdots + s_{i_k} = 0$ or 1 modulo 2. In a noisy planted instance, the system is evaluated on a random solution and independent noise is added, while

in a random instance, the right-hand side is uniformly random. Alekhnovich conjectured that the two are hard to distinguish when $k = 3$ and $m = O(n)$. We give a sample-efficient reduction from solving noisy planted k-LIN instances to distinguishing them from random instances. Suppose that m-equation, n-variable instances of the two types are efficiently distinguishable with advantage ϵ. Then, we show that $O(m \cdot (m/\epsilon)^{2/k})$-equation, n-variable noisy planted k-LIN instances are efficiently solvable with probability $\exp -\tilde{O}((m/\epsilon)^{6/k})$. The solver is based on a new approximate local list-decoding algorithm for the k-XOR code at large distances. The k-XOR encoding of a function $F\colon \Sigma \to \{-1, 1\}$ is its k-th tensor power $F^k(x_1, \ldots, x_k) = F(x_1) \cdots F(x_k)$. Given oracle access to a function G that μ-correlates with F^k, our algorithm outputs the description of a message that $(\mu^{1/k} - \epsilon)$-correlates with F with probability $\exp -\tilde{O}(k^2 \mu^{-2/k} \epsilon^{-2})$. Previous decoders have a worse dependence on μ (Levin, Combinatorica 1987) or do not apply to subconstant $\mu^{1/k}$.

Andrej Bogdanov
Chinese University of Hong Kong
Chinese University of Hong Kong
andrejb@cse.cuhk.edu.hk

Manuel Sabin, Prashant Nalini Vasudevan
University of California, Berkeley
msabin@berkeley.edu, prashantv91@gmail.com

CP16

Seeded Graph Matching via Large Neighborhood Statistics

We study a well known noisy model of the graph isomorphism problem. In this model, the goal is to perfectly recover the vertex correspondence between two edge-correlated graphs, with an initial seed set of correctly matched vertex pairs revealed as side information. For seeded problems, our result provides a dramatic improvement over previously known results. We show that it is possible to achieve the information-theoretic limit of graph sparsity in time polynomial in the number of vertices n. Moreover, we show the number of seeds needed for exact recovery in polynomial-time can be as low as n^ϵ in the sparse graph regime with the average degree smaller than n^ϵ) and $\Omega(\log n)$ in the dense graph regime. Our results also shed light on the unseeded problem. In particular, we give (the first) sub-exponential time algorithms for sparse models and an $n^{O(\log n)}$ algorithm for dense models for some parameters, including some that are not covered by recent results of Barak et al. Unlike previous work on graph matching, which used small neighborhoods or small subgraphs with a logarithmic number of vertices in order to match vertices, our algorithms match vertices if their large neighborhoods have a significant overlap in the number of seeds.

Elchanan Mossel
MIT
elmos@mit.edu

Jiaming Xu
Duke University
jiaming.xu868@duke.edu

CP17

Alternative Multicriteria Routes

We consider the problem of computing a set of alternative routes in a multicriteria setting where several network metrics are available. Previous approaches for alternative route computation were based on relaxing a single metric to obtain alternative routes whereas our approach for the multicriteria setting produces routes that are always optimal for a convex combination of the metrics. For the concrete example of route planning for bicycles with three natural metrics (distance, positive height difference, unsuitability for cycling) we show how to efficiently generate very natural alternative bicycle routes.

Florian Barth
University of Stuttgart
barth@fmi.uni-stuttgart.de

Stefan Funke
FMI, University of Stuttgart, Germany
funke@fmi.uni-stuttgart.de

Sabine Storandt
Universität Konstanz
sabine.storandt@uni-konstanz.de

CP17

Concatenated k-Path Covers

Given a directed graph $G(V, E)$, a k-(Shortest) Path Cover is a subset C of the nodes V such that every simple (or shortest) path in G consisting of k nodes contains at least one node from C. In this paper, we extend the notion of k-Path Covers such that the objects to be covered don't have to be single paths but can be concatenations of up to p simple (or shortest) paths. For the generalized problem of computing concatenated k-(Shortest) Path Covers, we present theoretical results regarding the VC-dimension of the concatenated path set in dependency of p as well as (approximation) algorithms. Subsequently, we study interesting special cases of concatenated k-Path Covers, in particular, covers for piecewise shortest paths, round tours and trees. For those, we show how the pruning algorithm for k-Path Cover computation can be abstracted and modified in order to also solve concatenated k-Path Cover problems. An extensive experimental study on different graph types proves the applicability and efficiency of our approaches.

Moritz Beck
Universität Konstanz
moritz.beck@uni-konstanz.de

Kam-Yiu Lam
City University of Hong Kong
cskylam@cityu.edu.hk

Joseph Kee Yin Ng
Hong Kong Baptist University
jng@comp.hkbu.edu.hk

Sabine Storandt
Universität Konstanz
sabine.storandt@uni-konstanz.de

Chun Jiang Zhu

University of Conneticut
chunjiang.zhu@uconn.edu

CP17

Fast and Exact Public Transit Routing with Restricted Pareto Sets

We present a novel exact journey planning approach to computing a reasonable subset of multi-criteria Pareto sets in public transit networks. Our restriction is well defined and independent of the choice of algorithm. In order to compute the restricted Pareto set efficiently, we present *Bounded McRAPTOR*, a new set of algorithms that extend the well-known McRAPTOR algorithm. The fastest variant employs a novel pruning scheme based on carefully computed bounds. Experiments on large metropolitan networks show that a four-criteria restricted Pareto set can be computed faster by a factor of up to 65, while retaining the important journeys of the full Pareto set. This easily enables interactive applications in practice, making multi-criteria Pareto-optimal journey planning scalable without the need of a preprocessing-based speedup technique.

Daniel Delling, Julian Dibbelt, Thomas Pajor
Apple Inc
ddelling@apple.com, jdibbelt@apple.com, tpajor@apple.com

CP17

Batch-Parallel Euler Tour Trees

The dynamic trees problem is to maintain a forest undergoing edge insertions and deletions while supporting queries for information such as connectivity. Though many data structures for this problem exist, few can exploit parallelism in the batch setting in which large batches of edges are inserted or deleted from the forest at once. We demonstrate that the Euler tour tree (ETT), an existing sequential dynamic trees data structure, can be parallelized in the batch setting. Our ETTs process a batch of k updates over an n-vertex forest with $O(k \log(1 + n/k))$ expected work and $O(\log n)$ depth with high probability. Our work bound is asymptotically optimal, and we improve on the depth bound achieved by Acar et al. for the batch-parallel dynamic trees problem. A crucial building block for our ETTs is a batch-parallel skip list data structure, which may be of independent interest. ETTs require a sequence data structure capable of joins and splits. We show that skip lists support batches of joins or splits of size k over n elements with $O(k \log(1 + n/k))$ expected work and $O(\log n)$ depth with high probability. We also achieve the same efficiency bounds for augmented skip lists, which lets us augment our ETTs to support subtree queries. Our data structures achieve between 67-96x self-relative speedup on 72 cores with hyper-threading on large batch sizes, and they also significantly outperform the fastest existing sequential alternatives empirically.

Thomas Tseng, Laxman Dhulipala
Carnegie Mellon University
thomasts@alumni.cmu.edu, ldhulipa@andrew.cmu.edu

Guy Blelloch
Computer Science Department
Carnegie Mellon University
guyb@cs.cmu.edu

CP18

A PTAS for Euclidean TSP with Hyperplane Neighborhoods

In the Traveling Salesperson Problem with Neighborhoods (TSPN), we are given a collection of geometric regions in some space. The goal is to output a tour of minimum length that visits at least one point in each region. Even in the Euclidean plane, TSPN is known to be APX-hard, which gives rise to studying more tractable special cases of the problem. In this paper, we focus on the fundamental special case of regions that are hyperplanes in the d-dimensional Euclidean space. While for $d = 2$ an exact algorithm with running time $O(n^5)$ is known, settling the exact approximability of the problem for $d = 3$ has been repeatedly posed as an open question. To date, only an approximation algorithm with guarantee exponential in d is known, and NP-hardness remains open. For arbitrary fixed d, we develop a Polynomial Time Approximation Scheme (PTAS) that works for both the tour and path version of the problem. Our algorithm is based on approximating the convex hull of the optimal tour by a convex polytope of bounded complexity. Such polytopes are represented as solutions of a sophisticated LP formulation, which we combine with the enumeration of crucial properties of the tour. As part of our analysis we develop a general sparsification technique to transform an arbitrary convex polytope into one with a constant number of vertices and, in turn, into one of bounded complexity in the above sense. Hereby, we maintain important properties of the polytope.

Antonios Antoniadis, Krzysztof Fleszar
Max-Planck-Institut für Informatik
aantonia@mpi-inf.mpg.de, kfleszar@mpi-inf.mpg.de

Ruben Hoeksma
Universität Bremen
hoeksma@uni-bremen.de

Kevin Schewior
Technische Universität München
École Normale Supérieure
kschewior@gmail.com

CP18

Polynomial-time Approximation Scheme for Minimum k-Cut in Planar and Minor-free Graphs

The k-cut problem asks, given a connected graph G and a positive integer k, to find a minimum-weight set of edges whose removal splits G into k connected components. We give the first polynomial-time algorithm with approximation factor $2 - \epsilon$ (with constant $\epsilon > 0$) for the k-cut problem in planar and minor-free graphs. Applying more complex techniques, we further improve our method and give a polynomial-time approximation scheme for the k-cut problem in both planar and minor-free graphs. Despite the persistent effort, to the best of our knowledge, this is the first improvement for the k-cut problem over standard approximation factor of 2 in any major class of graphs.

Alireza Farhadi
University of Maryland
farhadi94@gmail.com

MohammadHossein Bateni
Google Research

bateni@google.com

MohammadTaghi Hajiaghayi
University of Maryland, College Park
hajiagha@cs.umd.edu

CP18

Embedding Planar Graphs into Low-treewidth Graphs with Applications to Efficient Approximation Schemes for Metric Problems

Abstract not available.

Eli Fox-Epstein
Tufts University
ef@cs.tufts.edu

Philip Klein
Brown University
klein@brown.edu

Aaron Schild
University of California, Berkeley
aschild@berkeley.edu

CP18

Nearly ETH-tight Algorithms for Planar Steiner Tree with Terminals on Few Faces

The Steiner Tree problem is one of the most fundamental NP-complete problems as it models many network design problems. Recall that an instance of this problem consists of a graph with edge weights, and a subset of vertices (often called terminals); the goal is to find a subtree of the graph of minimum total weight that connects all terminals. A seminal paper by Erickson et al. [Math. Oper. Res., 1987] considers instances where the underlying graph is planar and all terminals can be covered by the boundary of k faces. Erickson et al. show that the problem can be solved by an algorithm using $n^{O(k)}$ time and $n^{O(k)}$ space, where n denotes the number of vertices of the input graph. In the past 30 years there has been no significant improvement of this algorithm, despite several efforts. In this work, we give an algorithm for Planar Steiner Tree with running time $2^{O(k)}n^{O(\sqrt{k})}$ using only polynomial space. Furthermore, we show the running time of our algorithm is almost tight: we prove that there is no $f(k)n^{o(\sqrt{k})}$ algorithm for Planar Steiner Tree for any computable function f, unless the Exponential Time Hypothesis fails.

Sándor Kisfaludi-Bak
Eindhoven University of Technology
s.kisfaludi.bak@tue.nl

Jesper Nederlof
TU Eindhoven
j.nederlof@tue.nl

Erik Jan van Leeuwen
Utrecht University
e.j.vanleeuwen@uu.nl

CP18

Contraction Decomposition in Unit Disk Graphs and Algorithmic Applications in Parameterized Complexity

We give a new decomposition theorem in unit disk graphs (UDGs) and demonstrate its applicability in the fields of Structural Graph Theory and Parameterized Complexity. First, we show that the class of UDGs admits a Contraction Decomposition Theorem. Prior studies on this topic exhibited that the classes of planar graphs [Klein, 2008], graphs of bounded genus [Demaine, Hajiaghayi and Mohar, 2010] and H-minor free graphs [Demaine, Hajiaghayi and Kawarabayashi, 2011] admit a Contraction Decomposition Theorem. Additionally, this result answers an open question posed by Hajiaghayi. Second, we present a "parameteric version' of our new decomposition theorem. We prove that there is an algorithm that given a UDG G and a positive integer k, runs in polynomial time and outputs a collection of $O(k)$ tree decompositions of G with the following properties. Each bag in any of these tree decompositions can be partitioned into $O(k)$ connected "pieces'. Moreover, for any subset S of at most k edges in G, there is a tree decomposition in the collection such that S is "well preserved' in the decomposition in the following sense. For any bag in the tree decomposition and any edge in S with both endpoints in the bag, either its endpoints lie in different pieces or they lie in a piece which is a clique. Having this decomposition at hand, we show that the design of parameterized algorithms for some cut problems becomes elementary.

Meirav Zehavi
Ben-Gurion University
zehavimeirav@gmail.com

Fahad Panolan
University of Bergen, Norway
fahad.panolan@ii.uib.no

Saket Saurabh
IMSc +UiB
saket@imsc.res.in

CP19

The Streaming k-Mismatch Problem

We consider the streaming complexity of a fundamental task in approximate pattern matching: the k-mismatch problem. In this problem, we must compute Hamming distances between a pattern of length n and all length-n substrings of a text for which the Hamming distance does not exceed a given threshold k. In our problem formulation, we report not only the Hamming distance but also, on demand, the full mismatch information, that is the list of mismatched pairs of symbols and their indices. The twin challenges of streaming pattern matching derive from the need both to achieve small working space and to guarantee that every arriving input symbol is processed quickly. We present a streaming algorithm for the k-mismatch problem which uses $O(k \log n \log \frac{n}{k})$ bits of space and spends $O(\log \frac{n}{k}(\sqrt{k \log k} + \log^3 n))$ time on each symbol of the input stream. In our formulation, the pattern also arrives in the stream, directly before the text. The running time almost matches the classic offline solution and the space usage is within a logarithmic factor of optimal. Our new algorithm therefore effectively resolves and extends a problem first introduced in FOCS'09. En route to this solution, we also give a deterministic $O(k(\log \frac{n}{k} + \log |\Sigma|))$-bit encoding of all the alignments with Hamming distance at most k of a length-n pattern within a text of length $O(n)$. This secondary result provides an optimal solution to a natural

encoding problem.

Raphael Clifford
University of Bristol, UK
raphael.clifford@bristol.ac.uk

Tomasz Kociumaka
Institute of Informatics
University of Warsaw
kociumaka@mimuw.edu.pl

Ely Porat
Bar Ilan University
porately@cs.biu.ac.il

CP19

Efficiently Approximating Edit Distance Between Pseudorandom Strings

We present an algorithm for approximating the edit distance $\mathrm{ed}(x, y)$ between two strings x and y in time parameterized by the degree to which one of the strings x satisfies a natural pseudorandomness property. The pseudorandomness model is asymmetric in that no requirements are placed on the second string y, which may be constructed by an adversary with full knowledge of x. We say that x is (p, B)-*pseudorandom* if all pairs a and b of disjoint B-letter substrings of x satisfy $\mathrm{ed}(a, b) \geq pB$. Given parameters p and B, our algorithm computes the edit distance between a (p, B)-pseudorandom string x and an arbitrary string y within a factor of $O(1/p)$ in time $\tilde{O}(nB)$, with high probability. If x is generated at random, then with high probability it will be $(\Omega(1), O(\log n))$-pseudorandom, allowing us to compute $\mathrm{ed}(x, y)$ within a constant factor in near linear time. For strings x of varying degrees of pseudorandomness, our algorithm offers a continuum of runtimes. Our algorithm is robust in the sense that it can handle a small portion of x being adversarial (i.e., not satisfying the pseudorandomness property). In this case, the algorithm incurs an additive approximation error proportional to the fraction of x which behaves maliciously.

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

CP19

Approximating Lcs in Linear Time: Beating the $\sqrt{n}$ Barrier

Longest common subsequence () is one of the most fundamental problems in combinatorial optimization. Apart from theoretical importance, has enormous applications in bioinformatics, revision control systems, and data comparison programs[1]. Although a simple dynamic program computes in quadratic time, it has been recently proven that the problem admits a conditional lower bound and may not be solved in truly subquadratic time [**?**]. In addition to this, is notoriously hard with respect to approximation algorithms. Apart from a trivial sampling technique that obtains a n^x approximation solution in time $O(n^{2-2x})$ nothing else is known for . This is in sharp contrast to its dual problem *edit distance* for which several linear time solutions are obtained in the past two decades [**?**, **?**, **?**, **?**, **?**]. Whether or not a nontrivial approximation for is possible in linear time has been raised as an open question by experts in the community [**?**]. In this work, we present the first nontrivial algorithm for approximating in linear time. Our main result is a linear time algorithm for the longest common subsequence which has an approximation factor of $O(n)$. This beats the $\sqrt{n}$ barrier for approximating in linear time.

[1] a notable example is the UNIX application diff

Saeed Seddighin, MohammadTaghi Hajiaghayi
University of Maryland, College Park
saeedreza.seddighin@gmail.com, hajiagha@cs.umd.edu

Masoud Seddighin
Sharif University of Technology
seddighin.masood@gmail.com

Xiaorui Sun
University of Illinois at Chicago
sunsirius@gmail.com

CP19

Lower Bounds for Text Indexing with Mismatches and Differences

In this paper we study lower bounds for the fundamental problem of text indexing with mismatches and differences. In this problem we are given a long string of length n, the "text', and the task is to preprocess it into a data structure such that given a query string Q, one can quickly identify substrings that are within Hamming or edit distance at most k from Q. This problem is at the core of various problems arising in biology and text processing. We start by demonstrating conditional lower bounds for $k = \Theta(\log n)$. We show that assuming the Strong Exponential Time Hypothesis, any data structure for text indexing that can be constructed in polynomial time cannot have $(n^{1-\delta})$ query time, for any $\delta > 0$. This bound also extends to the setting where we only ask for $(1 + \epsilon)$-approximate solutions for text indexing. However, in many applications the value of k is rather small, and one might hope that for small k we can develop more efficient solutions. We show that this would require a radically new approach as using the current methods one cannot avoid exponential dependency on k either in the space, or in the time bound for all even $\frac{8}{\sqrt{3}}\sqrt{\log n} \leq k = o(\log n)$. Our lower bounds also apply to the dictionary look-up problem, where instead of a text one is given a set of strings.

Tatiana Starikovskaya
Ecole Normale Supérieure
Paris
tat.starikovskaya@gmail.com

Vincent Cohen-Addad
Sorbonne Universités and CNRS
vcohenad@gmail.com

Laurent Feuilloley
IRIF, Univ Paris Diderot
feuilloley@irif.fr

CP19

Few Matches or Almost Periodicity: Faster Pattern Matching with Mismatches in Compressed Texts

A fundamental problem on strings in the realm of approximate string matching is *pattern matching with mismatches*: Given a text t, a pattern p, and a number k, determine whether some substring of t has Hamming distance at most k to p; such a substring is called a k-

match. We study the case of searching for a small pattern p in a text t that is compressed by a straight-line program. This grammar compression is popular in the string community as it unifies many well-known compression schemes such as the Lempel-Ziv family, dictionary methods, and others. We denote by m the length of p and by n the compressed size of t. While exact pattern matching, i.e., the case $k = 0$, is known to be solvable in near-linear time $\tilde{O}(n + m)$ [Jeż TALG'15], despite considerable interest in the string community, the fastest known algorithm for pattern matching with mismatches runs in time $\tilde{O}(n\sqrt{m}\,\mathrm{poly}(k))$ [Gawrychowski, Straszak ISAAC'13], which is far from linear even for very small k. In this paper, we obtain an algorithm for pattern matching with mismatches running in time $\tilde{O}((n+m)\,\mathrm{poly}(k))$. This is near-linear in the input size for any constant (or slightly superconstant) k. Our algorithm is based on a new structural insight for approximate pattern matching, essentially showing that either the number of k-matches is very small or both text and pattern must be almost periodic.

Karl Bringmann
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
kbringma@mpi-inf.mpg.de

Marvin Künnemann
Max-Planck-Institut für Informatik, Saarbrücken,
Germany
marvin@mpi-inf.mpg.de

Philip Wellnitz
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
wellnitz@mpi-inf.mpg.de

CP20

Polynomial-time Algorithm for Maximum Weight Independent Set on P_6-Free Graphs

In the classic Maximum Weight Independent Set problem we are given a graph G with a nonnegative weight function on vertices, and the goal is to find an independent set in G of maximum possible weight. While the problem is NP-hard in general, we give a polynomial-time algorithm working on any P_6-free graph, that is, a graph that has no path on 6 vertices as an induced subgraph. This improves the polynomial-time algorithm on P_5-free graphs of Lokshtanov et al., and the quasipolynomial-time algorithm on P_6-free graphs of Lokshtanov et al. The main technical contribution leading to our main result is enumeration of a polynomial-size family F of vertex subsets with the following property: for every maximal independent set I in the graph, F contains all maximal cliques of some minimal chordal completion of G that does not add any edge incident to a vertex of I.

Michal Pilipczuk
University of Warsaw
michal.pilipczuk@mimuw.edu.pl

Andrzej Grzesik
Jagiellonian University
Cracow, Poland
andrzej.grzesik@tcs.uj.edu.pl

Tereza Klimošová
Charles University
Prague, Czech Republic
tereza@kam.mff.cuni.cz

Marcin Pilipczuk
University of Warsaw
Warsaw, Poland
marcin.pilipczuk@mimuw.edu.pl

CP20

How to Guess an n-Digit Number

In a deductive game for two players, SF and PGOM, SF conceals an n-digit number $x = x_1, \ldots, x_n$ in base q, and PGOM, who knows n and q, tries to identify x by asking a number of questions, which are answered by SF. Each question is an n-digit number $y = y_1, \ldots, y_n$ in base q; each answer is the number of subscripts i such that $x_i = y_i$. Moreover, we require PGOM send all the questions at once. We show that the minimum number of questions required to determine x is $(2 + o_q(1))n/\log_q n$. Our result closes the gap between the lower bound attributed to Erdős and Rényi and the upper bounds developed subsequently by Lindström, Chvátal, Kabatianski, Lebedev and Thorpe. A more general problem is to determine the asymptotic formula of the metric dimension of Cartesian powers of a graph. We state the class of graphs for which the formula can be determined, and the smallest graphs for which we did not manage to settle.

Zilin Jiang
Technion - Israel Institute of Technology
zilinj@mit.edu

Nikita Polyanskii
Skolkovo Institute of Science and Technology
nikitapolyansky@gmail.com

CP20

The Maximum Number of Minimal Dominating Sets in a Tree

A tree with n vertices has at most $95^{n/13}$ minimal dominating sets. The growth constant $\lambda = \sqrt[13]{95} \approx 1.4194908$ is best possible. It is obtained as a kind of "dominant eigenvalue" of a bilinear operation on sixtuples that is derived from the dynamic-programming recursion for computing the number of minimal dominating sets of a tree. The semi-automatic computer-assisted way in which the growth constant was derived might be interesting in its own right. We also develop an output-sensitive algorithm for listing all minimal dominating sets with linear set-up time and linear delay between reporting successive solutions.

Günter Rote
Freie Universität Berlin
Institut für Informatik
rote@inf.fu-berlin.de

CP20

Vector Clique Decompositions

Let $\mathcal{F}_k$ be the set of k-vertex graphs. For a graph G, a k-decomposition is a set of induced subgraphs of G, each isomorphic to an element of $\mathcal{F}_k$, such that each pair of vertices of G is in exactly one element of the set. It is a fundamental result of Wilson that for all $n = |V(G)|$ sufficiently large, G has a k-decomposition if and only if G is k-divisible, namely $k - 1$ divides $n - 1$ and $\binom{k}{2}$ divides $\binom{n}{2}$. Let $\mathbf{v} \in \mathbf{R}^{|\mathcal{F}_k|}$ be indexed by $\mathcal{F}_k$. For a k-

decomposition L, let $\nu_{\mathbf{v}}(L) = \sum_{F \in \mathcal{F}_k} \mathbf{v}_F d_{L,F}$ where $d_{L,F}$ is the fraction of elements of L isomorphic to F. Let $\nu_{\mathbf{v}}(G) = \max_L \nu_{\mathbf{v}}(L)$, $\nu_{\mathbf{v}}(n) = \min\{\nu_{\mathbf{v}}(G) : |V(G)| = n\}$, $\nu_{\mathbf{v}} = \lim_{n\to\infty} \nu_{\mathbf{v}}(n)$. Replacing k-decompositions with their fractional relaxations, one obtains the fractional analogue $\nu^*_{\mathbf{v}}(G)$ and the corresponding $\nu^*_{\mathbf{v}}(n)$ and $\nu^*_{\mathbf{v}}$. Our main result is that for each $\mathbf{v} \in \mathbf{R}^{|\mathcal{F}_k|}$

$$\nu_{\mathbf{v}} = \nu^*_{\mathbf{v}} .$$

Furthermore, there is a polynomial time algorithm that produces a decomposition L of a k-decomposable graph such that $\nu_{\mathbf{v}}(L) \geq \nu_{\mathbf{v}} - o_n(1)$. Similar results hold in the directed and edge-colored settings. We use these results to obtain new and improved bounds on several decomposition results.

Raphael Yuster
Department of Mathematics
University of Haifa
raphy@math.haifa.ac.il

CP20

Four-Coloring P_6-Free Graphs

In this paper we present a polynomial time algorithm for the 4-coloring problem and the 4-precoloring extension problem restricted to the class of graphs with no induced six-vertex path, thus proving a conjecture of Huang. Combined with previously known results this completes the classification of the complexity of the 4-coloring problem for graphs with a connected forbidden induced subgraph.

Maria Chudnovsky, Sophie Spirkl
Princeton University
mchudnov@math.princeton.edu, sspirkl@princeton.edu

Mingxian Zhong
Columbia University
mz2325@columbia.edu

CP21

A New Integer Linear Program for the Steiner Tree Problem with Revenues, Budget and Hop Constraints

The Steiner tree problem with revenues, budgets and hop constraints (STPRBH) is a variant of the classical Steiner tree problem. This problem asks for a subtree in a given graph with maximum revenues corresponding to its nodes, where its total edge costs respect the given budget, and the number of edges between each node and its root does not exceed the hop limit. We introduce a new binary linear program with polynomial size based on partial ordering, which (up to our knowledge) for the first time solves all STPRBH instances from the DIMACS benchmark set to optimality. The set contains graphs with up to 500 nodes and 12500 edges.

Adalat Jabrayilov
TU Dortmund University
Department of Computer Science
adalat.jabrayilov@tu-dortmund.de

Petra Mutzel
Dortmund University
Germany
petra.mutzel@cs.uni-dortmund.de

CP21

Exactly Solving the Maximum Weight Independent Set Problem on Large Real-world Graphs

One powerful technique to solve NP-hard optimization problems in practice is branch-and-reduce search—which is branch-and-bound that intermixes branching with reductions to decrease the input size. While this technique is known to be very effective in practice for unweighted problems, very little is known for weighted problems, in part due to a lack of known effective reductions. In this work, we develop a full suite of new reductions for the maximum weight independent set problem and provide extensive experiments to show their effectiveness in practice on real-world graphs of up to millions of vertices and edges. Our experiments indicate that our approach is able to outperform existing state-of-the-art algorithms, solving many instances that were previously infeasible. In particular, we show that branch-and-reduce is able to solve a large number of instances up to two orders of magnitude faster than existing (inexact) local search algorithms—and is able to solve the majority of instances within 15 minutes. For those instances remaining infeasible, we show that combining kernelization with local search produces higher-quality solutions than local search alone.

Darren Strash
Department of Computer Science
Colgate University
darren.strash@gmail.com

Sebastian Lamm
Karlsruhe Institute of Technology
lamm@kit.edu

Christian Schulz
University of Vienna
Karlsruhe Institute of Technology
christian.schulz@univie.ac.at

Robert Williger
Karlsruhe Institute of Technology
willigerrobert@gmail.com

Huashuo Zhang
Colgate University
hzhang@colgate.edu

CP21

SAT-Encodings for Treecut Width and Treedepth

Graph decompositions associated with so-called width parameters (such as treewidth) have been the focus of extensive theoretical research. Finding an optimal decomposition is usually an NP-hard task. In this paper we propose, implement, and test the first practical decomposition algorithms for the width parameters treecut width and treedepth, which have recently gained a lot of attention in the theoretical research community. However, one prominent obstacle for any practical or experimental use of these two width parameters is the lack of any practical or implemented algorithm for actually computing the associated decompositions. Our approach for computing treecut width and treedepth decompositions is based on efficient encodings into the propositional satisfiability problem (SAT). Once an encoding is generated, any satisfiability

solver can be used to find the decomposition. This allows us to leverage the surprising power of todays state-of-the art SAT solvers. The success of SAT-based decomposition methods crucially depends on the used characterisation of the decomposition method, as not every characterisation is suitable for that task. For treecut width and treedepth, we propose new characterisations that are based on sequences of partitions of the vertex set, a method that was pioneered for clique-width. We implemented and systematically tested our encodings on various benchmark instances, including famous named graphs and random graphs of various density.

Robert Ganian
TU Wien
rganian@gmail.com

Neha Lodha
TU Wien, Vienna, Austria
nehalodha73@gmail.com

Sebastian Ordyniak
University of Sheffield, Sheffield, UK
sordyniak@gmail.com

Stefan Szeider
TU Wien
stefan@szeider.net

Pan Peng
University of Sheffield
p.peng@sheffield.ac.uk

CP21

Efficiently Enumerating Hitting Sets of Hypergraphs Arising in Data Profiling

We devise an enumeration method for inclusion-wise minimal hitting sets in hypergraphs. It has delay $O(m^{k^*+1}n^2)$ and uses linear space. Hereby, n is the number of vertices, m the number of hyperedges, and k^* the rank of the transversal hypergraph. In particular, on classes of hypergraphs for which the cardinality k^* of the largest minimal hitting set is bounded, the delay is polynomial. The algorithm solves the extension problem for minimal hitting sets as a subroutine. We show that the extension problem is W[3]-complete when parameterised by the cardinality of the set which is to be extended. For the subroutine, we give an algorithm that is optimal under the exponential time hypothesis. Despite these lower bounds, we provide empirical evidence showing that the enumeration outperforms the theoretical worst-case guarantee on hypergraphs arising in the profiling of relational databases, namely, in the detection of unique column combinations.

Thomas Bläsius, Tobias Friedrich, Julius Lischeid
Hasso Plattner Institute
University of Potsdam
thomas.blaesius@hpi.de, tobias.friedrich@hpi.de,
julius.lischeid@student.hpi.de

Kitty Meeks
University of Glasgow
kitty.meeks@glasgow.ac.uk

Martin Schirneck
Hasso Plattner Institute
University of Potsdam
martin.schirneck@hpi.de

CP22

A Faster External Memory Priority Queue with DecreaseKeys

A priority queue is a fundamental data structure that maintains a dynamic set of (key, priority)-pairs and supports Insert, Delete, ExtractMin and DecreaseKey operations. In the external memory model, the current best priority queue supports each operation in amortized $O(\frac{1}{B}\log\frac{N}{B})$ I/Os. If the DecreaseKey operation does not need to be supported, one can design a more efficient data structure that supports the Insert, Delete and ExtractMin operations in $O(\frac{1}{B}\log\frac{N}{B}/\log\frac{M}{B})$ I/Os. A recent result shows that a degradation in performance is inevitable by proving a lower bound of $\Omega(\frac{1}{B}\log B/\log\log N)$ I/Os for priority queues with DecreaseKeys. In this paper we tighten the gap between the lower bound and the upper bound by proposing a new priority queue which supports the DecreaseKey operation and has an expected amortized I/O complexity of $O(\frac{1}{B}\log\frac{N}{B}/\log\log N)$. Our result improves the external memory priority queue with DecreaseKeys for the first time in over a decade, and also gives the fastest external memory single source shortest path algorithm.

Shunhua Jiang
Institute for Interdisciplinary Information Sciences,
Tsinghua University
gingerbell00@gmail.com

Kasper G. Larsen
MADALGO, Aarhus University
larsen@cs.au.dk

CP22

Optimal Construction of Compressed Indexes for Highly Repetitive Texts

We propose algorithms that, given the input string of length n over integer alphabet of size σ, construct the Burrows–Wheeler transform (BWT), the permuted longest-common-prefix (PLCP) array, and the LZ77 parsing in $O(n/\log_\sigma n + r\,\mathrm{polylog}\,n)$ time and working space, where r is the number of runs in the BWT of the input. These are the essential components of many compressed indexes such as compressed suffix tree, FM-index, and grammar and LZ77-based indexes, but also find numerous applications in sequence analysis and data compression. The value of r is a common measure of repetitiveness that is significantly smaller than n if the string is highly repetitive. Since just accessing every symbol of the string requires $\Omega(n/\log_\sigma n)$ time, the presented algorithms are time and space optimal for inputs satisfying the assumption $n/r \in \Omega(\mathrm{polylog}\,n)$ on the repetitiveness. For such inputs our result improves upon the currently fastest general algorithms of Belazzougui (STOC 2014) and Munro et al. (SODA 2017) which run in $O(n)$ time and use $O(n/\log_\sigma n)$ working space. We also show how to use our techniques to obtain optimal solutions on highly repetitive data for other fundamental string processing problems such as: Lyndon factorization, construction of run-length compressed suffix arrays, and some classical "textbook' problems such as computing the longest substring occurring at least some fixed number of times.

Dominik Kempa
Department of Computer Science

University of Helsinki
dominik.kempa@gmail.com

CP22

Strategies for Stable Merge Sorting

We introduce new stable natural merge sort algorithms, called 2-merge sort and α-merge sort. We prove upper and lower bounds for several merge sort algorithms, including Timsort, Shiver's sort, α-stack sorts, and our new 2-merge and α-merge sorts. The upper and lower bounds have the forms $c \cdot n \log m$ and $c \cdot n \log n$ for inputs of length n comprising m runs. For Timsort, we prove a lower bound of $(1.5 - o(1))n \log n$. For 2-merge sort, we prove optimal upper and lower bounds of approximately $(1.089 \pm o(1))n \log m$. We prove similar asymptotically matching upper and lower bounds for α-merge sort, when $\varphi < \alpha < 2$, where φ is the golden ratio. Our bounds are in terms of merge cost; this upper bounds the number of comparisons and accurately models runtime. The merge strategies can be used for any stable merge sort, not just natural merge sorts. The new 2-merge and α-merge sorts have better worst-case merge cost upper bounds and are slightly simpler to implement than the widely-used Timsort; they also perform better in experiments.

Samuel Buss
Deaprtment of Mathematics
UC San Diego
sbuss@ucsd.edu

Alexander Knop
Department of Mathematics
UC San Diego
aknop@ucsd.edu

CP22

A New Path from Splay to Dynamic Optimality

Consider the task of performing a sequence of searches in a binary search tree. After each search, an algorithm is allowed to arbitrarily restructure the tree, at a cost proportional to the amount of restructuring performed. The cost of an algorithm's execution is the sum of the time spent searching and the time spent optimizing those searches with restructuring operations. This notion was introduced by Sleator and Tarjan in 1985, along with an algorithm and a conjecture. The algorithm, Splay, is an elegant procedure for performing adjustments while moving searched items to the top of the tree. The conjecture, called "dynamic optimality,' is that the cost of splaying is always within a constant factor of the optimal algorithm for performing searches. The conjecture stands to this day. We offer the first systematic proposal for settling the dynamic optimality conjecture. At the heart of our methods is what we term a simulation embedding: a formula that maps executions to lists of keys which induce a target algorithm to simulate the execution. We build a simulation embedding for Splay by inducing it to perform arbitrary subtree transformations, and use this to show that if the cost of splaying a sequence of items upper bounds the cost of splaying every subsequence thereof, then splay is dynamically optimal. We call this the subsequence property. Building on this machinery, we further show that the subsequence property is also a necessary condition for dynamic optimality.

Caleb Levy
Princeton University
cclevy@princeton.edu

Robert Tarjan
Princeton University
Microsoft Research
ret@cs.princeton.edu

CP22

A Sort of an Adversary

We describe an *efficient* deterministic adversary that forces any comparison-based sorting algorithm to perform at least $\frac{37}{64}\, n \log n$ comparisons. This improves on previous efficient adversaries of Atallah and Kosaraju (1981), Richards and Vaidya (1988), and of Bordal et al. (1996) that force any sorting algorithm to perform at least $\frac{1}{2} \cdot n \log n$ comparisons.

Or Zamir
Tel Aviv University
orzamir@mail.tau.ac.il

Haim Kaplan
Tel-Aviv University
haimk@post.tau.ac.il

Uri Zwick
Tel Aviv University
zwick@tau.ac.il

CP23

Iterative Refinement for ℓ_p-Norm Regression

We give improved algorithms for the ℓ_p-regression problem, $\min_x ||x||_p$ such that $Ax = b$, for all $p \in (1, 2) \cup (2, \infty)$. Our algorithms obtain a high accuracy solution in $\tilde{O}_p(m^{\frac{|p-2|}{2p+|p-2|}}) \leq \tilde{O}_p(m^{1/3})$ iterations, where each iteration requires solving a linear system. Incorporating a procedure for maintaining an approximate inverse of the linear systems that we need to solve at each iteration, we give an algorithm for solving ℓ_p-regression to $1/poly(n)$ accuracy that runs in time $\tilde{O}_p(m^{\max\{\omega, 7/3\}})$, where ω is the matrix multiplication constant. For the current best value of ω, this means that we can solve ℓ_p regression for all constant p bounded away from 1 as fast as ℓ_2 regression. The iteration counts, as well as running times on general matrices and sparse graphs improve up previous best result by [Bubeck-Cohen-Lee-Li STOC'18], as well as other, more general purpose convex optimization algorithms. At the core of our algorithms is an iterative refinement scheme for ℓ_p-norms, using the quadratically-smoothed ℓ_p-norms introduced in the work of Bubeck et al. Formally, we specify a minimization problem over the quadratically-smoothed ℓ_p norms, such that a crude solution to this problem allows us to improve the quality of an approximate ℓ_p-norm minimizer by constant factor, leading to algorithms with fast convergence.

Sushant Sachdeva, Deeksha Adil
University of Toronto
sachdeva@cs.toronto.edu, deeksha@cs.toronto.edu

Rasmus Kyng
Yale University
rjkyng@gmail.com

Richard Peng
Georgia Tech

rpeng@cc.gatech.edu

CP23

Perron-Frobenius Theory in Nearly Linear Time: Positive Eigenvectors, M-Matrices, Graph Kernels, and Other Applications

In this paper we provide nearly linear time algorithms for several problems closely associated with the classic Perron-Frobenius theorem, including computing Perron vectors, i.e. entrywise non-negative eigenvectors of non-negative matrices, and solving linear systems in asymmetric M-matrices, a generalization of Laplacian systems. The running times of our algorithms depend nearly linearly on the input size and polylogarithmically on the desired accuracy and problem condition number. Leveraging these results we also provide improved running times for a broader range of problems including computing random walk-based graph kernels, computing Katz centrality, and more. The running times of our algorithms improve upon previously known results which either depended polynomially on the condition number of the problem, required quadratic time, or only applied to special cases. We obtain these results by providing new iterative methods for reducing these problems to solving linear systems in Row-Column Diagonally Dominant (RCDD) matrices. Our methods are related to the classic shift-and-invert preconditioning technique for eigenvector computation and constitute the first alternative to the result in Cohen et al. (2016) for reducing stationary distribution computation and solving directed Laplacian systems to solving RCDD systems.

Amirmahdi Ahmadinejad, Arun Jambulapati
Stanford University
ahmadi@stanford.edu, jmblpati@stanford.edu

Amin Saberi
Management Science and Engineering
Stanford University
saberi@stanford.edu

Aaron Sidford
Stanford University
sidford@stanford.edu

CP23

Approximability of P -¿ Q Matrix Norms: Generalized Krivine Rounding and Hypercontractive Hardness

We study the problem of computing the $p \to q$ norm of a matrix A in R^{mxn}, defined as $||A||_{p\to q} := \max_{x\in R^n\{0\}} ||Ax||_q/||x||_p$. This problem generalizes the spectral norm of a matrix ($p = q = 2$) and the Grothendieck problem (p=infinity, q=1), and has been widely studied in various regimes. When $p >= q$, the problem exhibits a dichotomy: constant factor approximation algorithms are known if 2 is in [q,p], and the problem is hard to approximate within almost polynomial factors when 2 is not in [q,p]. For the case when 2 is in [q,p] we prove almost matching approximation and NP-hardness results. The regime when $p < q$, known as hypercontractive norms, is particularly significant for various applications but much less well understood. The case with $p = 2$ and $q > 2$ was studied by [Barak et. al., STOC'12] who gave sub-exponential algorithms for a promise version of the problem (which captures small-set expansion) and also proved hardness of approximation results based on the Exponential Time Hypothesis. However, no NP-hardness of approximation is known for these problems for any $p < q$. We prove the first NP-hardness result for approximating hypercontractive norms. We show that for any $1 < p < q <$ infinity with 2 not in $[p, q], ||A||_{p\to q}$ is hard to approximate within $2^{O((\log n)^{1-\epsilon})}$ assuming NP is not contained in BPTIME$2^{(\log n)^{O(1)}}$.

Vijay Bhattiprolu
CMU
vpb@cs.cmu.edu

Mrinalkanti Ghosh
Toyota Technological Institute Chicago
mkghosh@ttic.edu

Venkatesan Guruswami, Euiwoong Lee
Carnegie Mellon University
guruswami@cmu.edu, euiwoong@cims.nyu.edu

Madhur Tulsiani
TTI Chicago
madhurt@ttic.edu

CP23

Optimizing Quantum Optimization Algorithms via Faster Quantum Gradient Computation

We develop a quantum algorithm that computes the gradient of a multi-variate real-valued function $f : \mathbb{R}^d \to \mathbb{R}$ by evaluating it only a logarithmic number of times in superposition. Our algorithm is an improved version of Jordan's gradient computation algorithm, providing an approximation of the gradient ∇f with quadratically better dependence on the evaluation accuracy of f, for an important class of smooth functions. Furthermore, we show that most objective functions arising during the training of variational quantum circuits satisfy the necessary smoothness conditions, hence our algorithm improves the complexity of computing their gradient. We also show that in a continuous phase-query model, our gradient computation algorithm has optimal query complexity up to poly-logarithmic factors, for a class of smooth functions. Moreover, we show that for low-degree multivariate polynomials our algorithm can provide exponential speedups compared to Jordan's algorithm in terms of the dimension d. We provide *efficient* subroutines for performing a delicate interconversion between probability and phase oracles incurring only a logarithmic overhead, which might be of independent interest. Finally, using these tools we improve the runtime of prior approaches for training quantum auto-encoders, variational quantum eigensolvers (VQE), and quantum approximate optimization algorithms (QAOA).

Andras Gilyen
QuSoft, CWI and University of Amsterdam
gilyen@cwi.nl

Srinivasan Arunachalam
Center of Theoretical Physics, MIT
QuSoft / CWI
arunacha@mit.edu

Nathan Wiebe
Station Q QuArC, Microsoft Research

nawiebe@microsoft.com

CP23

Proportional Volume Sampling and Approximation Algorithms for A-Optimal Design

We study the optimal design problems where the goal is to choose a set of linear measurements to obtain the most accurate estimate of an unknown vector in d dimensions. We study the A-optimal design variant where the objective is to minimize the average variance of the error in the maximum likelihood estimate of the vector being measured. The problem also finds applications in sensor placement in wireless networks, sparse least squares regression, feature selection for k-means clustering, and matrix approximation. In this paper, we introduce proportional volume sampling to obtain improved approximation algorithms for A-optimal design. Our main result is to obtain improved approximation algorithms for the A-optimal design problem by introducing the proportional volume sampling algorithm. Our results nearly optimal bounds in the asymptotic regime when the number of measurements done, k, is significantly more than the dimension d. We also give first approximation algorithms when k is small including when $k = d$. The proportional volume-sampling algorithm also gives approximation algorithms for other optimal design objectives such as D-optimal design and generalized ratio objective matching or improving previous best known results. Interestingly, we show that a similar guarantee cannot be obtained for the E-optimal design problem. We also show that the A-optimal design problem is NP-hard to approximate within a fixed constant when $k = d$.

Uthaipon Tantipongpipat
Georgia Institute of Technology
Computer Science
uthaipon@gmail.com

Mohit Singh
H. Milton Stewart School of Industrial & Systems Engineering
Georgia Institute of Technology
mohitsinghr@gmail.com

Aleksandar Nikolov
University of Toronto
anikolov@cs.toronto.edu

CP24

Polynomial Planar Directed Grid Theorem

The grid theorem [Robertson, Seymour 1986] is a central result in the study of graph minors and has found many algorithmic applications. The relation between treewidth and grid minors is polynomial, even linear, for planar graphs [Robertson, Seymour and Thomas, 1994] enabling many important consequences (e.g. sublinear separators and subexponential algorithms). In the mid-90s, Reed and Johnson, Robertson, Seymour and Thomas proposed a notion of directed treewidth. They conjectured an excluded grid theorem for directed graphs, which was proved in 2015 by the latter two authors but the function relating directed treewidth and grid minors is big, even in the planar case. Directed grids have found algorithmic applications such as low-congestion routing. However, in the undirected case the polynomial bound on the size of grid minors in planar graphs have made this tool so successful. The lack of such a bound has so far prevented further applications in the directed setting. The main result of this paper is to establish a polynomial bound for the directed grid theorem on planar digraphs. We think this will enable further applications of directed treewidth and directed grids. We also give a "treewidth sparsifier' for directed graphs, which already was considered in undirected graphs. This allows us to obtain an Eulerian subgraph of bounded degree that still has high directed treewidth. We believe this result is of independent interest for structure graph theory.

Meike Hatzel
Technische Universität Berlin
meike.hatzel@tu-berlin.de

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

Stephan Kreutzer
Technical University Berlin
stephan.kreutzer@tu-berlin.de

CP24

A Tight Erdős-Pósa Function for Planar Minors

Let H be a planar graph. By a classical result of Robertson and Seymour, there is a function $f :\to$ such that for all $k \in$ and all graphs G, either G contains k vertex-disjoint subgraphs each containing H as a minor, or there is a subset X of at most $f(k)$ vertices such that $G - X$ has no H-minor. We prove that this remains true with $f(k) = ck \log k$ for some constant $c = c(H)$. This bound is best possible, up to the value of c, and improves upon a recent result of Chekuri and Chuzhoy [STOC 2013], who established this with $f(k) = ck \log^d k$ for some universal constant d. The proof is constructive and yields a polynomial-time $O(\log)$-approximation algorithm for packing subgraphs containing an H-minor.

Jean-Florent Raymond
TU Berlin
raymond@tu-berlin.de

Wouter Cames Van Batenburg
Département d'Informatique, Université Libre de Bruxelles
Brussels, Belgium
wcamesva@ulb.ac.be

Tony Huynh
Université Libre de Bruxelles
tony.bourbaki@gmail.com

Gwenaël Joret
Département d'Informatique, Université Libre de Bruxelles
Brussels, Belgium
gjoret@ulb.ac.be

CP24

Polynomial Bounds for Centered Colorings on Proper Minor-Closed Graph Classes

Centered colorings play an important role in the theory of sparse graph classes introduced by Nešetřil and Ossona de Mendez, as they structurally characterize classes of bounded expansion: a class $\mathcal{C}$ has bounded expansion if and only if there is a function $f\colon \mathbb{N} \to \mathbb{N}$ such that every graph $G \in \mathcal{C}$ for every $p \in \mathbb{N}$ admits a p-centered coloring

with at most $f(p)$ colors. Unfortunately, known proofs of the existence of such colorings yield large upper bounds on the function f governing the number of colors needed, even for as simple classes as planar graphs. In this paper, we prove that every K_t-minor-free graph admits a p-centered coloring with $O(p^{g(t)})$ colors for some function g. This provides the first polynomial upper bounds on the number of colors needed in p-centered colorings of graphs drawn from proper minor-closed classes, which answers an open problem posed by Dvořák. As an algorithmic application, we use our main result to prove that if $\mathcal{C}$ is a fixed proper minor-closed class of graphs, then given graphs H and G, on p and n vertices, respectively, where $G \in \mathcal{C}$, it can be decided whether H is a subgraph of G in time $2^{O(p \log p)} \cdot n^{O(1)}$ and space $n^{O(1)}$.

Sebastian Siebertz
Technische Universität Berlin
siebertz@informatik.hu-berlin.de

Michal Pilipczuk
University of Warsaw
michal.pilipczuk@mimuw.edu.pl

Marcin Pilipczuk
University of Warsaw
Warsaw, Poland
marcin.pilipczuk@mimuw.edu.pl

CP24

Every Collinear Set in a Planar Graph is Free

We show that if a planar graph G has a plane straight-line drawing in which a subset S of its vertices is collinear, then for any set of points X with $|X| = |S|$, there is a plane straight-line drawing of G in which the vertices in S are mapped to the points in X. This solves an open problem posed by Ravsky and Verbitsky in 2008. In their terminology, we show that every collinear set is free. This result has applications in graph drawing, including untangling, column planarity, universal point subsets, and partial simultaneous drawings.

Vida Dujmovic
University of Ottawa
vida.dujmovic@uottawa.ca

Fabrizio Frati
Università Roma Tre
frati@dia.uniroma3.it

Daniel Gonçalves
LIRMM (CNRS & Université de Montpellier)
goncalves@lirmm.fr

Pat Morin
Carleton University
morin@scs.carleton.ca

Günter Rote
Freie Universität Berlin
Institut für Informatik
rote@inf.fu-berlin.de

CP24

Towards Tight(er) Bounds for the Excluded Grid Theorem

We study the Excluded Grid Theorem, a fundamental structural result in graph theory, that was proved by Robertson and Seymour in their seminal work on graph minors. The theorem states that there is a function $f : \mathbb{Z}^+ \to \mathbb{Z}^+$, such that for every integer $g > 0$, every graph of treewidth at least $f(g)$ contains the $(g \times g)$-grid as a minor. For every integer $g > 0$, let $f(g)$ be the smallest value for which the theorem holds. Establishing tight bounds on $f(g)$ is an important graph-theoretic question. Robertson and Seymour showed that $f(g) = \Omega(g^2 \log g)$ must hold. For a long time, the best known upper bounds on $f(g)$ were super-exponential in g. The first polynomial upper bound of $f(g) = O(g^{98} \log g)$ was proved by Chekuri and Chuzhoy. It was later improved to $f(g) = O(g^{36} \log g)$, and then to $f(g) = O(g^{19} \log g)$. In this paper we further improve this bound to $f(g) = O(g^9 \log g)$. We believe that our proof is significantly simpler than the proofs of the previous bounds. Moreover, while there are natural barriers that seem to prevent the previous methods from yielding tight bounds for the theorem, it seems conceivable that the techniques proposed in this paper can lead to even tighter bounds on $f(g)$.

Julia Chuzhoy
Toyota Technological Institute at Chicago
cjulia@ttic.edu

Zihan Tan
The University of Chicago
zihantan@uchicago.edu

CP25

A Practical Algorithm for Spatial Agglomerative Clustering

We study an agglomerative clustering problem motivated by visualizing disjoint glyphs centered at specific locations on a geographic map. As we zoom out, the glyphs grow and start to overlap. We replace overlapping glyphs by one larger merged glyph to maintain disjointness. Our goal is to compute the resulting hierarchical clustering efficiently in practice. A straightforward algorithm for such spatial agglomerative clustering runs in $O(n^2 \log n)$ time, where n is the number of glyphs. This is not efficient enough for many real-world datasets which contain up to tens or hundreds of thousands of glyphs. Recently the theoretical upper bound was improved to $O(n\alpha(n) \log^7 n)$ time [**?**], where $\alpha(n)$ is the inverse Ackermann function. Although this new algorithm is asymptotically much faster than the naive algorithm, from a practical point of view, it does not perform better for $n \leq 10^6$. In this paper we present a new agglomerative clustering algorithm which works efficiently in practice. Our algorithm relies on the use of quadtrees to speed up spatial computations. Interestingly, even in non-pathological datasets we can encounter large glyphs that intersect many quadtree cells and that are involved in many clustering events. We therefore devise a special strategy to handle such large glyphs. We test our algorithm on several synthetic and real-world datasets and show that it performs well in practice.

Thom Castermans
TU Eindhoven
t.h.a.castermans@tue.nl

Bettina Speckmann
Dept. of Mathematics and Computer Science
TU Eindhoven

b.speckmann@tue.nl

Kevin Verbeek
TU Eindhoven
k.a.b.verbeek@tue.nl

CP25

Practical Methods for Computing Large Covering Tours and Cycle Covers with Turn Cost

We study the problem of computing provably optimal and near-optimal solutions for the NP-hard problem of finding covering tours and cycle covers with turn cost, which are of practical importance for a variety of applications, such as pest control and precision farming. Previous work has largely focused on theoretical aspects, such as complexity and approximation. We develop a number of algorithm engineering techniques and refinements to make such theoretical insights practically useful, resulting in a comprehensive study for solving a wide spectrum of large instances. We compute provably optimal solutions for instances with more than 1000 pixels, from the largest previous solved instance size of 76 (de Assis and de Souza 2011). Making use of additional algorithm engineering techniques for handling very large instances, we also compute near-optimal solutions for instances with up to 300000 pixels, for which we give solutions that are typically within a few percent of our computed lower bounds. We also provide an experimental comparison of a practically refined version of our new theoretical approach with the approximation technique of Arkin et al. that dates back to 2001; we show that our new LP/IP-based approximation method closes 70% of the remaining optimality gap to the lower bound.

Sándor Fekete, Dominik M. Krupke
TU Braunschweig
s.fekete@tu-bs.de, krupke@ibr.cs.tu-bs.de

CP25

Faster Support Vector Machines

The time complexity of support vector machines (SVMs) prohibits training on huge data sets with millions of samples. Recently, multilevel approaches to train SVMs have been developed to allow for time efficient training on huge data sets. While regular SVMs perform the entire training in one - time consuming - optimization step, multilevel SVMs first build a hierarchy of problems decreasing in size that resemble the original problem and then train an SVM model for each hierarchy level benefiting from the solved models of previous levels. We present a faster multilevel support vector machine that uses a label propagation algorithm to construct the problem hierarchy. Extensive experiments show that our new algorithm achieves speed-ups of up to two orders of magnitude while having similar or better classification quality over state-of-the-art algorithms.

Sebastian Schlag
Karlsruhe Institute of Technology
Institute for Theoretical Informatics, Algorithmics II
sebastian.schlag@kit.edu

Matthias Schmitt
Karlsruhe Institute of Technology
uydvo@student.kit.edu

Christian Schulz
University of Vienna
Karlsruhe Institute of Technology
christian.schulz@univie.ac.at

CP25

Scalable Edge Partitioning

Edge-centric distributed computations have appeared as a recent technique to improve the shortcomings of think-like-a-vertex algorithms on large scale-free networks. In order to increase parallelism on this model, *edge partitioning*—partitioning edges into roughly equally sized blocks—has emerged as an alternative to traditional (node-based) graph partitioning. In this work, we develop a fast parallel split-and-connect graph construction algorithm in the distributed setting and show that combining our parallel construction with advanced parallel node partitioning algorithms yields high-quality edge partitions in a scalable way. Our technique scales to networks with billions of edges, and runs efficiently on thousands of PEs. Our extensive experiments show that our algorithm computes solutions of high quality on large real-world networks and large hyperbolic random graphs—which have a power law degree distribution and are therefore specifically targeted by edge partitioning.

Darren Strash
Department of Computer Science
Colgate University
darren.strash@gmail.com

Sebastian Schlag
Karlsruhe Institute of Technology
Institute for Theoretical Informatics, Algorithmics II
sebastian.schlag@kit.edu

Christian Schulz
University of Vienna
Karlsruhe Institute of Technology
christian.schulz@univie.ac.at

Daniel Seemaier
Karlsruhe Institute of Technology
daniel.seemaier@student.kit.edu

CP25

Parallel Range, Segment and Rectangle Queries with Augmented Maps

The support of range, segment and rectangle queries are fundamental problems in computational geometry, and have extensive applications in many domains. Despite significant theoretical work on these problems, efficient implementations can be complicated, and most implementations do not have useful theoretical bounds. In this paper, we focus on simple and efficient parallel algorithms and implementations for range, segment and rectangle queries, which have worst-case bounds in theory and good performance in practice. Our approach uses an abstract data type called augmented map, based on which we develop both multilevel tree structures and sweepline algorithms supporting range, segment and rectangle queries. For the sweepline algorithms, we propose a parallel paradigm and show corresponding cost bounds. Theoretically, the construction algorithms of all of our data structures are work-efficient and highly parallelized. We implemented all the data structures in the paper, ten in all, using a parallel augmented map library. Based on the library, each data structure only requires about 100 lines of C++ code. We test their per-

formance on large data sets and a machine with 72-cores (144 hyperthreads). Our implementation achieves 32-68x speedup in construction, and up to 126x in queries. Sequentially, our implementation outperforms or is competitive to existing libraries including the CGAL library and the Boost library.

Yihan Sun
Carnegie Mellon University
yihans@cs.cmu.edu

Guy Blelloch
Computer Science Department
Carnegie Mellon University
guyb@cs.cmu.edu

CP26

Lift and Project Algorithms for Precedence Constrained Scheduling to Minimize Completion Time

We consider the classic problem of scheduling jobs with precedence constraints on a set of identical machines to minimize the weighted completion time objective. Understanding the exact approximability of the problem when job lengths are uniform is a well known open problem in scheduling theory. In this paper, we show an optimal algorithm that runs in polynomial time and achieves an approximation factor of (2+epsilon) for the weighted completion time objective when the number of machines is a constant. The result is obtained by building on the lift and project approach introduced in a breakthrough work by Levey and Rothvoss for the makespan minimization problem.

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

Shashwat Garg
TU Eindhoven
garg.shashwat@gmail.com

Shi Li
Toyota Technological Institute at Chicago
shil@buffalo.edu

CP26

A Polynomial Time Constant Approximation For Minimizing Total Weighted Flow-time

Abstract not available.

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

Uriel Feige
Weizmann Institute, Israel
uriel.feige@weizmann.ac.il

Shi Li
Toyota Technological Institute at Chicago
shil@buffalo.edu

CP26

A New Dynamic Programming Approach for Spanning Trees with Chain Constraints and Beyond

Short spanning trees subject to additional constraints capture many interesting problem settings and are important building blocks in various approximation algorithms. We consider spanning trees subject to constraints on the edges in a family of cuts forming a laminar family of small width. Our main contribution is a new dynamic programming approach, where the value of a table entry does not only depend on the values of previous table entries, as is usually the case, but also on a specific representative solution saved together with each table entry. This allows for handling a broad range of constraint types. In combination with other techniques—including negatively correlated rounding and a polyhedral approach—we obtain several new results. We first present a quasi-polynomial time algorithm for the Minimum Chain-Constrained Spanning Tree problem with an essentially optimal guarantee: violation of chain constraints is bounded by a 1+epsilon factor, and the cost is no larger than that of an optimal solution not violating any chain constraint. The best previous procedure is a bicriteria approximation violating each chain constraint by up to a constant factor and losing another factor in the objective. Moreover, our approach can naturally handle lower bounds on the chain constraints. Furthermore, we show how our approach can also handle parity constraints as used in the context of (path) TSP and a generalization thereof, and discuss implications in this context.

Martin Nagele, Rico Zenklusen
ETH Zurich
martin.naegele@ifor.math.ethz.ch, ricoz@math.ethz.ch

CP26

On Approximating (Sparse) Covering Integer Programs

We consider approximation algorithms for covering integer programs parametrized by column sparsity. First, we show that a simple algorithm based on randomized rounding with alteration improves or matches the best known approximation algorithms in a wide range of parameter settings, and these bounds are essentially optimal. As a byproduct of the simplicity of the alteration algorithm and analysis, we can derandomize the algorithm without any loss in the approximation guarantee or efficiency. Non-trivial approximation algorithms for covering integer programs are based on solving the natural LP relaxation strengthened with knapsack cover inequalities. Our second contribution is a fast (essentially near-linear time) approximation scheme for solving the strengthened LP with a factor of n speed up over the previous best running time.

Chandra Chekuri, Kent Quanrud
University of Illinois at Urbana-Champaign
chekuri@illinois.edu, quanrud2@illinois.edu

CP26

A 1.5-Approximation for Path TSP

We present a 1.5-approximation for the Metric Path Traveling Salesman Problem (Path TSP). All recent improvements on Path TSP crucially exploit a structural property shown by An, Kleinberg, and Shmoys [Journal of the ACM, 2015], namely that narrow cuts with respect to a Held-Karp solution form a chain. We significantly deviate from these approaches by showing the benefit of dealing with larger s-t cuts, even though they are much less structured. More precisely, we show that a variation of the dynamic programming idea recently introduced by Traub and Vygen [SODA, 2018] is versatile enough to deal with larger size cuts, by exploiting a seminal result of Karger

on the number of near-minimum cuts. This avoids a recursive application of dynamic programming as used by Traub and Vygen, and leads to a considerably simpler algorithm avoiding an additional error term in the approximation guarantee. We match the still unbeaten 1.5-approximation guarantee of Christofides' algorithm for TSP. Hence, any further progress on the approximability of Path TSP will also lead to an improvement for TSP. Speaker: Martin Nagele, ETH Zrich, Switzerland

Rico Zenklusen
ETH Zurich
ricoz@math.ethz.ch

CP27

Coresets Meet EDCS: Algorithms for Matching and Vertex Cover on Massive Graphs

We present a single unified approach for designing improved algorithms for matching and vertex cover across several models of computation for processing massive graphs such as streaming, distributed communication, and the massively parallel computation (MPC) models. For example, we give:

- The first one pass, significantly-better-than-2-approximation for matching in the random arrival order streaming model that uses subquadratic space, namely a 1.5-approximation streaming algorithm that uses $\tilde{O}(n^{1.5})$ space.
- The first 2 round, better-than-2-approximation for matching in the MPC model that uses subquadratic space per machine, namely a 1.5-approximation algorithm with $\tilde{O}(\sqrt{mn}+n)$ memory per machine.

By building on our unified approach, we further develop parallel algorithms in the MPC model that give a $(1+\epsilon)$-approximation to matching and an $O(1)$-approximation to vertex cover in only $O(\log\log n)$ MPC rounds and $O(n/\text{polylog}(n))$ memory per machine. These results settle multiple open questions posed by Czumaj et.al [STOC 2018]. We obtain our results by a novel combination of two previously disjoint set of techniques, namely randomized composable coresets and edge degree constrained subgraphs (EDCS). We significantly extend the power of these techniques and prove several new structural results.

Sepehr Assadi
University of Pennsylvania
sassadi@cis.upenn.edu

MohammadHossein Bateni
Google Research
bateni@google.com

Aaron Bernstein
Technical University of Berlin, Germany
bernstei@gmail.com

Vahab Mirrokni
Google Inc.
mirrokni@google.com

Cliff Stein
Columbia University
cliff@ieor.columbia.edu

CP27

Sparsifying Distributed Algorithms with Ramifications in Massively Parallel Computation and Centralized Local Computation

Abstract not available.

Mohsen Ghaffari, Jara Uitto
ETH Zurich
ghaffari@inf.ethz.ch, jara.uitto@inf.ethz.ch

CP27

Low Congestion Cycle Covers and their Applications

A cycle cover of a bridgeless graph G is a collection of simple cycles in G such that each edge e appears on at least one cycle. Motivated by applications to distributed computation, we introduce the notion of *low-congestion* cycle covers, in which all cycles in the cycle collection are both *short* and nearly *edge-disjoint*. Formally, a $(,)$-cycle cover of a graph G is a collection of cycles in G in which each cycle is of length at most and each edge participates in at least one cycle and at most cycles. Perhaps quite surprisingly, we prove the following: Every bridgeless graph of diameter D admits a $(,)$-cycle cover where $=\tilde{O}(D)$ and $=\tilde{O}(1)$. That is, the edges of G can be covered by cycles such that each cycle is of length at most $\widetilde{O}(D)$ and each edge participates in at most $\widetilde{O}(1)$ cycles. These parameters are existentially tight up to polylogarithmic terms. We demonstrate the usefulness of low congestion cycle covers in different settings of resilient computation. For instance, we consider a Byzantine fault model where in each round, the adversary chooses a single message and corrupt in an arbitrarily manner. We provide a compiler that turns any r-round distributed algorithm for a graph G with diameter D, into an equivalent fault tolerant algorithm with $r\cdot(D)$ rounds.

Merav Parter, Eylon Yogev
Weizmann Institute
merav,parter@weizmann.ac.il, eylony@gmail.com

CP27

Massively Parallel Approximation Algorithms for Edit Distance and Longest Common Subsequence

String similarity measures are among the most fundamental problems in computer science. The notable examples are *edit distance* () and *longest common subsequence* (). These problems find their applications in various contexts such as computational biology, text processing, compiler optimization, data analysis, image analysis, etc. In this work, we revisit edit distance and longest common subsequence in the parallel settings. We present massively parallel algorithms for both problems that are optimal in the following senses:

- The approximation factor of our algorithms is $1+\epsilon$.
- The round complexity of our algorithms is constant.
- The total running time of our algorithms over all machines is (n^2). This matches the running time of the best-known solutions for approximating edit distance and longest common subsequence within a $1+\epsilon$ factor in the sequential setting.

Our main technical contribution is a novel parallel algorithm for computing a set of compositions, and recursively decomposing each function into a set of smaller iterative compositions (in terms of memory needed to solve the problem). These two methods together give us a strong tool for

approximating combinatorial problems. For instance, can be formulated as a recursive composition of functions and therefore this tool enables us to approximate within a factor $1 + \epsilon$.

Saeed Seddighin
University of Maryland, College Park
saeedreza.seddighin@gmail.com

Xiaorui Sun
University of Illinois at Chicago
sunsirius@gmail.com

MohammadTaghi Hajiaghayi
University of Maryland, College Park
hajiagha@cs.umd.edu

CP27

Distributed Algorithms Made Secure: A Graph Theoretic Approach

An inherent property of most existing distributed algorithms is that throughout the course of their execution, the nodes get to learn not only their own output but rather learn quite a lot on the outputs of many other entities. This leakage of information might be a major obstacle in many settings. In this paper, we introduce a new framework for *secure distributed graph algorithms* and provide the first *general compiler* that takes any "natural' non-secure distributed algorithm that runs in r rounds, and turns it into a secure algorithm that runs in $\widetilde{O}(r \cdot D \cdot (\Delta))$ rounds where Δ is the maximum degree in the graph and D is its diameter. The security of the compiled algorithm is information-theoretic but holds only against a semi-honest adversary that controls a single node in the network. This compiler is made possible due to a new combinatorial structure called *private neighborhood trees*: a collection of n trees $T(u_1), \ldots, T(u_n)$, one for each vertex $u_i \in V(G)$, such that each tree $T(u_i)$ spans the neighbors of u_i *without going through* u_i. In a $(,)$-private neighborhood trees each tree $T(u_i)$ has depth at most and each edge $e \in G$ appears in at most different trees. We show a construction where $= \widetilde{O}(\Delta \cdot D)$ and $= \widetilde{O}(D)$.

Merav Parter, Eylon Yogev
Weizmann Institute
merav,parter@weizmann.ac.il, eylony@gmail.com

CP28

Interval Vertex Deletion Admits a Polynomial Kernel

Given a graph G and an integer k, the Interval Vertex Deletion (IVD) problem asks whether there exists a subset $S \subseteq V(G)$ of size at most k such that $G - S$ is an interval graph. This problem is known to be NP-complete [Yannakakis, STOC'78]. Originally in 2012, Cao and Marx showed that IVD is fixed parameter tractable: they exhibited an algorithm with running time $10^k n^{\mathcal{O}(1)}$ [Cao and Marx, SODA'14]. The existence of a polynomial kernel for IVD remained a well-known open problem in Parameterized Complexity. In this paper, we settle this problem in the affirmative.

Akanksha Agrawal
Hungarian Academy of Sciences
akanksha.agrawal.2029@gmail.com

CP28

Quantum Speedups for Exponential-Time Dynamic Programming Algorithms

In this paper we study quantum algorithms for NP-complete problems whose best classical algorithm is an exponential time application of dynamic programming. We introduce the path in the hypercube problem that models many of these dynamic programming algorithms. In this problem we are asked whether there is a path from 0^n to 1^n in a given subgraph of the Boolean hypercube, where the edges are all directed from smaller to larger Hamming weight. We give a quantum algorithm that solves path in the hypercube in time $O^*(1.817^n)$. The technique combines Grover's search with computing a partial dynamic programming table. We use this approach to solve a variety of vertex ordering problems on graphs in the same time $O^*(1.817^n)$, and graph bandwidth in time $O^*(2.946^n)$. Then we use similar ideas to solve the travelling salesman problem and minimum set cover in time $O^*(1.728^n)$.

Jevgenijs Vihrovs, Andris Ambainis, Kaspars Balodis, Janis Iraids, Martins Kokainis, Krisjanis Prusis
University of Latvia
jevgenijs.vihrovs@lu.lv, andris.ambainis@lu.lv, kbalodis@gmail.com, janis.iraids@gmail.com, martins.kokainis@gmail.com, krisjanis.prusis@lu.lv

CP28

A Time and Space-Optimal Algorithm for the Many-Visits TSP

The many-visits traveling salesperson problem (MV-TSP) asks for an optimal tour of n cities that visits each city c a prescribed number k_c of times. Travel costs may be asymmetric, and visiting a city twice in a row may incur a non-zero cost. The MV-TSP problem finds applications in scheduling, geometric approximation, and Hamiltonicity of certain graph families. The fastest known algorithm for MV-TSP is due to Cosmadakis and Papadimitriou (SICOMP, 1984). It runs in time $n^{O(n)} + O(n^3 \log \sum_c k_c)$ and requires $n^{\Omega(n)}$ space. The algorithm has a *logarithmic* dependence on the total length $\sum_c k_c$ of the tour, allowing it to handle instances with very long tours, beyond what is tractable in the standard TSP setting. In this paper we significantly improve on the said algorithm, giving an MV-TSP algorithm that runs in *single-exponential* time with *polynomial* space. More precisely, we obtain the run time $2^{O(n)} + O(n^3 \log \sum_c k_c)$, with $O(n^2 \log \sum_c k_c)$ space. The space requirement of our algorithm is (essentially) the size of the output, and assuming the Exponential-time Hypothesis (ETH), the time requirement is optimal. Our algorithm is deterministic, and arguably both simpler and easier to analyse than the original approach of Cosmadakis and Papadimitriou. It involves an optimization over directed spanning trees and a recursive, centroid-based decomposition of trees.

André Berger
Maastricht University
Department of Quantitative Economics
a.berger@maastrichtuniversity.nl

László Kozma
Eindhoven University of Technology
Department of Mathematics and Computer Science

l.kozma@tue.nl

Matthias Mnich
Maastricht University / Universit{
Department of Quantitative Economics / Computer Science
mmnich@uni-bonn.de

Roland Vincze
Maastricht University
r.vincze@maastrichtuniversity.nl

CP28

Losing Treewidth by Separating Subsets

We study the problem of deleting the smallest set S of vertices (resp. edges) from a given graph G such that the induced subgraph (resp. subgraph) $G \setminus S$ belongs to some class $\mathcal{H}$. We consider the case where graphs in $\mathcal{H}$ have treewidth bounded by t, and give a general framework to obtain approximation algorithms for both vertex and edge-deletion settings from approximation algorithms for certain natural graph partitioning problems called k-Subset Vertex Separator and k-Subset Edge Separator, respectively. For the vertex deletion setting, our framework combined with the current best result for k-Subset Vertex Separator, improves approximation ratios for basic problems such as k-Treewidth Vertex Deletion and Planar-$\mathcal{F}$ Vertex Deletion. Our algorithms are simpler than previous works and give the first deterministic and uniform approximation algorithms under the natural parameterization. For the edge deletion setting, we give improved approximation algorithms for k-Subset Edge Separator combining ideas from LP relaxations and important separators. We present their applications in bounded-degree graphs, and also give an APX-hardness result for the edge deletion problems.

Euiwoong Lee, Anupam Gupta, Jason M. Li
Carnegie Mellon University
euiwoong@cims.nyu.edu, anupamg@cs.cmu.edu, jmli@cs.cmu.edu

Pasin Manurangsi
UC Berkeley
pasin@berkeley.edu

Michal Wlodarczyk
University of Warsaw
m.wlodarczyk@minuw.edu.pl

CP28

On r-Simple k-Path and Related Problems Parameterized by k/r

In r-Simple k-Path, given a digraph G on n vertices and positive integers r, k, the goal is to decide whether G has an r-simple k-path, which is a walk where every vertex occurs at most r times and the total number of vertex occurrences is k. Abasi et al. (2014) obtained a randomized algorithm of running time $4^{(k/r)\log r} \cdot n^{O(1)}$ for this problem and a related problem called (r, k)-Monomial Detection. Gabizon et al. (2015) designed a deterministic $2^{O((k/r)\log r)} \cdot n^{O(1)}$-time algorithm for these two problems and a related problem called p-Set (r, q)-Packing. These results prove that the three problems are *single-exponentially* fixed-parameter tractable (FPT) when parameterized by the product of *two* parameters, that is, k/r and $\log r$. We consider the question from a wider perspective: are the above problems FPT when parameterized by k/r only? We resolve the wider question by (a) obtaining a $2^{O((k/r)^2 \log(k/r))} \cdot (n + \log k)^{O(1)}$-time algorithm for r-Simple k-Path on digraphs and a $2^{O(k/r)} \cdot (n + \log k)^{O(1)}$-time algorithm for r-Simple k-Path on undirected graphs, (b) showing that p-Set (r, q)-Packing is FPT (in contrast, we prove that p-Multiset (r, q)-Packing is W[1]-hard), and (c) proving that (r, k)-Monomial Detection is para-NP-hard even if only two distinct variables are in polynomial P and the circuit is non-canceling. All our algorithms are deterministic.

Meirav Zehavi
Ben-Gurion University
zehavimeirav@gmail.com

Gregory Gutin, Magnus Wahlstrom
Royal Holloway, University of London
g.gutin@rhul.ac.uk, magnus.wahlstrom@rhul.ac.uk

CP29

An Illuminating Algorithm for the Light Bulb Problem

The Light Bulb Problem is one of the most basic problems in data analysis. One is given as input n vectors in $\{-1, 1\}^d$, which are all independently and uniformly random, except for a planted pair of vectors with inner product at least $\rho \cdot d$ for some constant $\rho > 0$. The task is to find the planted pair. The most straightforward algorithm leads to a runtime of $\Omega(n^2)$. Algorithms based on techniques like Locality-Sensitive Hashing achieve runtimes of $n^{2-O(\rho)}$; as ρ gets small, these approach quadratic. Building on prior work, we give a new algorithm for this problem which runs in time $O(n^{1.582} + nd)$, *regardless* of how small ρ is. This matches the best known runtime due to Karppa et al. Our algorithm combines techniques from previous work on the Light Bulb Problem with the so-called 'polynomial method in algorithm design,' and has a simpler analysis than previous work. Our algorithm is also easily derandomized, leading to a deterministic algorithm for the Light Bulb Problem with the same runtime of $O(n^{1.582} + nd)$, improving previous results.

Josh Alman
MIT
jalman@mit.edu

CP29

A Framework for Searching in Graphs in the Presence of Errors

We consider a problem of searching for an unknown target vertex t in a graph. Each vertex-query points to a vertex v and the response either admits that v is the target or provides any neighbor s that lies on a shortest path from v to t. This model has been introduced for trees by Onak and Parys [FOCS 2006] and for general graphs by Emamjomeh-Zadeh et al. [STOC 2016]. In the latter, the authors provide an algorithm for the independent noise model where each query independently receives an erroneous answer with probability $p < 1/2$. We show an algorithm that needs at most $\frac{\log_2 n}{1-H(r)}$ queries under adversarial errors with error rate bounded by a constant $r < 1/2$. We then show that our algorithm coupled with a Chernoff bound argument leads to a simpler algorithm for the independent noise model and has query complexity that is both simpler and asymptotically better than the one of Emamjomeh-Zadeh et al. [STOC 2016]. Our approach has

a wide range of applications. First, it improves and simplifies the Robust Interactive Learning framework proposed by Emamjomeh-Zadeh and Kempe [NIPS 2017]. Secondly, analogous analysis for edge-queries, where query to edge e returns its endpoint that is closer to target, recovers an asymptotically optimal noisy binary search algorithm, matching the complexity of Feige et al. [SIAM J. Comput. 1994]. Thirdly, we improve and simplify upon an algorithm for searching in unbounded domains due to Aslam and Dhagat [STOC 1991].

Dariusz Dereniowski
Gdansk University of Technology
deren@eti.pg.gda.pl

Stefan Tiegel
Department of Computer Science
ETH Zurich
tiegels@student.ethz.ch

Przemyslaw Uznanski
ETH Zurich
przemyslaw.uznanski@inf.ethz.ch

Daniel Wolleb-Graf
Department of Computer Science
ETH Zurich
daniel.graf@inf.ethz.ch

CP29

Simple Concurrent Labeling Algorithms for Connected Components

We present some new concurrent labeling algorithms for finding connected components and study their theoretical efficiency. Even though many such algorithms have been proposed and many experiments with them have been done, our algorithms are simpler. We obtain an $O(\lg n)$ step bound for two of our algorithms using a novel multi-round analysis. We conjecture that our other algorithms also take $O(\lg n)$ steps but are unable to fully analyze them. We also point out some gaps in previous analyses of similar algorithms. Our results show that even a basic problem like connected components still has secrets to reveal.

Sixue Liu
Princeton University
sixuel@princeton.edu

Robert Tarjan
Princeton University
Microsoft Research
ret@cs.princeton.edu

CP29

Isotonic Regression by Dynamic Programming

For a given sequence of n numbers, we want to find a monotonically increasing sequence of the same length that best approximates it in the sense of minimizing the weighted sum of absolute values of the differences. A conceptually easy dynamic programming approach leads to an algorithm with running time $O(n \log n)$. While other algorithms with the same running time are known, our algorithm is very simple. The only auxiliary data structure that it requires is a priority queue. The approach extends to other error measures such as sum of squares.

Günter Rote
Freie Universität Berlin
Institut für Informatik
rote@inf.fu-berlin.de

CP29

Selection from Heaps, Row-Sorted Matrices and X+Y Using Soft Heaps

We use *soft heaps* to obtain simpler optimal algorithms for selecting the k-th smallest item, and the set of k smallest items, from a heap-ordered tree, from a collection of sorted lists, and from $X + Y$, where X and Y are two unsorted sets. Our results match, and in some ways extend and improve, classical results of Frederickson (1993) and Frederickson and Johnson (1982). In particular, for selecting the k-th smallest item, or the set of k smallest items, from a collection of m sorted lists we obtain a new optimal *"output-sensitive"* algorithm that performs only $O(m+\sum_{i=1}^{m} \log(k_i+1))$ comparisons, where k_i is the number of items of the i-th list that belong to the overall set of k smallest items.

Uri Zwick
Tel Aviv University
zwick@tau.ac.il

Haim Kaplan
Tel-Aviv University
haimk@post.tau.ac.il

Or Zamir
Tel Aviv University
orzamir@mail.tau.ac.il

László Kozma
Eindhoven University of Technology
Department of Mathematics and Computer Science
l.kozma@tue.nl

CP30

Communication-Rounds Tradeoffs for Common Randomness and Secret Key Generation

We study the role of interaction in the Common Randomness Generation (CRG) and Secret Key Generation (SKG) problems. In the CRG problem, two players, Alice and Bob, respectively get samples $X_1, X_2, \ldots$ and $Y_1, Y_2, \ldots$ with the pairs (X_1, Y_1), (X_2, Y_2), ... being drawn independently from some known probability distribution μ. They wish to communicate so as to agree on L bits of randomness. The SKG problem is the restriction of the CRG problem to the case where the key is required to be close to random even to an eavesdropper who can listen to their communication (but does not have access to the inputs of Alice and Bob). In this work, we study the relationship between the amount of communication and the number of rounds of interaction in both the CRG and the SKG problems. Specifically, we construct a family of distributions $\mu = \mu_{r,n,L}$, parametrized by integers r, n and L, such that for every r there exists a constant $b = b(r)$ for which CRG (respectively SKG) is feasible when $(X_i, Y_i) \sim \mu_{r,n,L}$ with $r+1$ rounds of communication, each consisting of $O(\log n)$ bits, but when restricted to $r/2 - 2$ rounds of interaction, the total communication must exceed $\Omega(n/\log^b(n))$ bits. Prior to our work no separations were known for $r \geq 2$.

Mitali Bafna
John A. Paulson School of Engineering and Applied

Sciences
Harvard University
mitalibafna@g.harvard.edu

Badih Ghazi
Google Research
badihghazi@gmail.com

Noah Golowich
Harvard University
ngolowich@college.harvard.edu

Madhu Sudan
Harvard
madhu@cs.harvard.edu

CP30

The Andoni-Krauthgamer-Razenshteyn Characterization of Sketchable Norms Fails for Sketchable Metrics

Andoni, Krauthgamer and Razenshteyn (AKR) proved that a finite-dimensional normed space $(X, \|\cdot\|_X)$ admits a $O(1)$ sketching algorithm (namely, with $O(1)$ sketch size and $O(1)$ approximation) if and only if for every $\epsilon \in (0,1)$ there exist $\alpha \geq 1$ and an embedding $f : X \to \ell_{1-\epsilon}$ such that $\|x-y\|_X \leq \|f(x)-f(y)\|_{1-\epsilon} \leq \alpha\|x-y\|_X$ for all $x, y \in X$. The "if part" of this theorem follows from a sketching algorithm of Indyk. The contribution of AKR is therefore to demonstrate that the mere availability of a sketching algorithm implies the existence of the aforementioned geometric realization. Indyk's algorithm shows that the "if part" of the AKR characterization holds true for any metric space whatsoever, i.e., the existence of an embedding as above implies sketchability even when X is not a normed space. Due to this, a natural question that AKR posed was whether the assumption that the underlying space is a normed space is needed for their characterization of sketchability. We resolve this question by proving that for arbitrarily large $n \in \mathbb{N}$ there is an n-point metric space $(M(n), d_{M(n)})$ which is $O(1)$-sketchable yet for every $\epsilon \in (0, \frac{1}{2})$, if $\alpha(n) \geq 1$ and $f_n : M(n) \to \ell_{1-\epsilon}$ are such that $d_{M(n)}(x,y) \leq \|f_n(x) - f_n(y)\|_{1-\epsilon} \leq \alpha(n) d_{M(n)}(x,y)$ for all $x, y \in M(n)$, then necessarily $\lim_{n\to\infty} \alpha(n) = \infty$.

Assaf Naor
Princeton University
naor@math.princeton.edu

Subhash Khot
New York University
khot@cs.nyu.edu

CP30

Tight Bounds for ℓ_p Oblivious Subspace Embeddings

An ℓ_p oblivious subspace embedding is a distribution over $r \times n$ matrices Π such that for any fixed $n \times d$ matrix A,

$$\Pr_{\Pi}[\text{for all } x,\ \|Ax\|_p \leq \|\Pi Ax\|_p \leq \kappa\|Ax\|_p] \geq 9/10,$$

where r is the *dimension* of the embedding, κ is the *distortion* of the embedding, and for an n-dimensional vector y, $\|y\|_p = \left(\sum_{i=1}^n |y_i|\right)^{1/p}$ is the ℓ_p-norm. Another important property is the *sparsity* of Π, that is, the maximum number of non-zero entries per column, as this determines the running time of computing $\Pi \cdot A$. While for $p = 2$ there are nearly optimal tradeoffs in terms of the dimension, distortion, and sparsity, for the important case of $1 \leq p < 2$, much less was known. In this paper we obtain nearly optimal tradeoffs for ℓ_p oblivious subspace embeddings for every $1 \leq p < 2$. Oblivious subspace embeddings are crucial for distributed and streaming environments, as well as entrywise ℓ_p low rank approximation. Our results give improved algorithms for these applications.

Ruosong Wang, David Woodruff
Carnegie Mellon University
ruosongw@andrew.cmu.edu, dwoodruf@cs.cmu.edu

CP30

Optimal Las Vegas Approximate Near Neighbors in ℓ_p

We show that approximate near neighbor search in high dimensions can be solved in a Las Vegas fashion (i.e., without false negatives) for ℓ_p $(1 \leq p \leq 2)$ while matching the performance of optimal locality-sensitive hashing. Specifically, we construct a data-independent Las Vegas data structure with query time $O(dn^\rho)$ and space usage $O(dn^{1+\rho})$ for (r, cr)-approximate near neighbors in $\mathbf{R}^d$ under the ℓ_p norm, where $\rho = 1/c^p + o(1)$. Furthermore, we give a Las Vegas locality-sensitive filter construction for the unit sphere that can be used with the data-dependent data structure of Andoni et al. (SODA 2017) to achieve optimal space-time tradeoffs in the data-dependent setting. For the symmetric case, this gives us a data-dependent Las Vegas data structure with query time $O(dn^\rho)$ and space usage $O(dn^{1+\rho})$ for (r, cr)-approximate near neighbors in $\mathbf{R}^d$ under the ℓ_p norm, where $\rho = 1/(2c^p - 1) + o(1)$.

Alexander Wei
Harvard University
weia@college.harvard.edu

CP30

Optimal Lower Bounds for Distributed and Streaming Spanning Forest Computation

We show optimal lower bounds for spanning forest computation in two different models: * One wants a data structure for fully dynamic spanning forest in which updates can insert or delete edges amongst a base set of n vertices. The sole allowed query asks for a spanning forest, which the data structure should successfully answer with some given (potentially small) constant probability $\epsilon > 0$. We prove that any such data structure must use $\Omega(n \log^3 n)$ bits of memory. * There is a referee and n vertices in a network sharing public randomness, and each vertex knows only its neighborhood; the referee receives no input. The vertices each send a message to the referee who then computes a spanning forest of the graph with constant probability $\epsilon > 0$. We prove the average message length must be $\Omega(\log^3 n)$ bits. Both our lower bounds are optimal, with matching upper bounds provided by the AGM sketch [AGM12] (which even succeeds with probability 1-1/poly(n)). Furthermore, for the first setting we show optimal lower bounds even for low failure probability δ, as long as $\delta > 2^{-n^{1-\epsilon}}$.

Huacheng Yu
Stanford University
yuhch123@gmail.com

Jelani Nelson
Harvard

minilek@seas.harvard.edu

CP31

A Deamortization Approach for Dynamic Spanner and Dynamic Maximal Matching

Many dynamic graph algorithms have an amortized update time, rather than a stronger worst-case guarantee. But amortized data structures are not suitable for real-time systems, where each individual operation has to be executed quickly. For this reason, there exist many recent randomized results that aim to provide a guarantee stronger than amortized expected. The strongest possible guarantee for a randomized algorithm is that it is always correct (Las Vegas), and has *high-probability worst-case* update time, which gives a bound on the time for each individual operation that holds with high probability. In this paper we present the first polylogarithmic high-probability worst-case time bounds for the dynamic spanner and the dynamic maximal matching problem. 1. For dynamic spanner, the only known $o(n)$ worst-case bounds were $O(n^{3/4})$ high-probability worst-case update time for maintaining a 3-spanner and $O(n^{5/9})$ for maintaining a 5-spanner. We give a $O(1)^k \log^3 n$ high-probability worst-case time bound for maintaining a $(2k-1)$-spanner, which yields the first worst-case polylog update time for all constant k. 2. For dynamic maximal matching, or dynamic 2-approximate maximum matching, no algorithm with $o(n)$ worst-case time bound was known and we present an algorithm with $O(\log^5 n)$ high-probability worst-case time; similar worst-case bounds existed only for maintaining a matching that was $(2+\epsilon)$-approximate, and hence not maximal.

Aaron Bernstein
Technical University of Berlin, Germany
bernstei@gmail.com

Sebastian Forster
University of Salzburg
forster@cs.sbg.ac.at

Henzinger Monika
University of Vienna
mhenzinger@gmail.com

CP31

Deterministically Maintaining a $(2+\epsilon)$-Approximate Minimum Vertex Cover in $O(1/\epsilon^2)$ Amortized Update Time

We consider the problem of maintaining an (approximately) minimum vertex cover in an n-node graph $G = (V, E)$ that is getting updated dynamically via a sequence of edge insertions/deletions. We show how to maintain a $(2+\epsilon)$-approximate minimum vertex cover, *deterministically*, in this setting in $O(1/\epsilon^2)$ amortized update time. Prior to our work, the best known deterministic algorithm for maintaining a $(2+\epsilon)$-approximate minimum vertex cover was due to Bhattacharya, Henzinger and Italiano [SODA 2015]. Their algorithm has an update time of $O(\log n/\epsilon^2)$. Recently, Bhattacharya, Chakrabarty, Henzinger [IPCO 2017] and Gupta, Krishnaswamy, Kumar, Panigrahi [STOC 2017] showed how to maintain an $O(1)$-approximation in $O(1)$-amortized update time for the same problem. Our result gives an *exponential* improvement over the update time of Bhattacharya et al. [SODA 2015], and nearly matches the performance of the *randomized* algorithm of Solomon [FOCS 2016] who gets an approximation ratio of 2 and an expected amortized update time of $O(1)$. We derive our result by analyzing, via a novel technique, a variant of the algorithm by Bhattacharya et al. Specifically, we consider an idealized setting where the update time of an algorithm can take any arbitrary fractional value, and use insights from this setting to come up with an appropriate potential function.

Sayan Bhattacharya
University of Warwick
Coventry
S.Bhattacharya@warwick.ac.uk

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

CP31

(1 + Eps)-Approximate Incremental Matching in Constant Deterministic Amortized Time

We study the matching problem in the incremental setting, where we are given a sequence of edge insertions and aim at maintaining a near-maximum cardinality matching of the graph with small update time. We present a deterministic algorithm that, for any constant $\varepsilon > 0$, maintains a $(1+\varepsilon)$-approximate matching with constant amortized update time per insertion.

Chris Schwiegelshohn
Sapienza University of Rome
schwiegelshohn@diag.uniroma1.it

Fabrizio Grandoni
IDSIA
fabrizio.grandoni@gmail.com

Stefano Leonardi
Sapienza University of Rome
leonardi@dis.uniroma1.it

Piotr Sankowski
Institute of Informatics, University of Warsaw
sank@mimuw.edu.pl

Shay Solomon
Tel Aviv University
solo.shay@gmail.com

CP31

Fully Dynamic Maximal Independent Set with Sublinear in n Update Time

The first fully dynamic algorithm for maintaining a maximal independent set (MIS) with update time that is sublinear in the number of edges was presented recently by the authors of this paper [Assadi et.al., STOC'18]. The algorithm is deterministic and its update time is $O(m^{3/4})$, where m is the number of edges. Subsequently, Gupta and Khan and independently Du and Zhang [arXiv, April 2018] presented deterministic algorithms for dynamic MIS with update times of $\tilde{O}(m^{2/3})$. Du and Zhang also gave a randomized algorithm with update time $\tilde{O}(\sqrt{m})$. Moreover, they provided some partial (conditional) hardness results hinting that the update time of $m^{1/2-\epsilon}$, and in particular $n^{1-\epsilon}$ for n-vertex dense graphs, is a natural barrier for this problem for any constant $\epsilon > 0$, for both deterministic and

randomized algorithms that satisfy a certain property. In this paper, we break this natural barrier and present the first fully dynamic (randomized) algorithm for maintaining an MIS with update time that is always sublinear in the number of vertices, namely, an $\tilde{O}(\sqrt{n})$ expected amortized update time. We also show that a simpler variant of our algorithm can achieve an $\tilde{O}(m^{1/3})$ expected amortized update time, which results in an improved performance over our $\tilde{O}(\sqrt{n})$ update time algorithm for sufficiently sparse graphs, and breaks the $m^{1/2}$ barrier for all values of m.

Sepehr Assadi
University of Pennsylvania
sassadi@cis.upenn.edu

Krzysztof Onak
IBM Research
konak@us.ibm.com

Baruch Schieber
New Jersey Institute of Technology
baruch.m.schieber@njit.edu

Shay Solomon
Tel Aviv University
solo.shay@gmail.com

CP31

Dynamic Edge Coloring with Improved Approximation

Given an undirected simple graph $G = (V, E)$ that undergoes edge insertions and deletions, we wish to efficiently maintain an edge coloring with only a few colors. The previous best dynamic algorithm by [**?**] could deterministically maintain a valid edge coloring using $2\Delta-1$ colors with $O(\log \Delta)$ update time, where Δ stands for the current maximum vertex degree of graph G. In this paper, we first propose a new static $(1+\epsilon)\Delta$ edge coloring algorithm that runs in near-linear time. Based on this static algorithm, we show that there is a randomized dynamic algorithm for this problem that only uses $(1+\epsilon)\Delta$ colors with $O(\log^8 n/\epsilon^4)$ amortized update time when $\Delta \geq \Omega(\log^2 n/\epsilon^2)$, where $\epsilon > 0$ is an arbitrarily small constant.

Ran Duan
Tsinghua University, Beijing, China
duanran@mail.tsinghua.edu.cn

Haoqing He, Tianyi Zhang
IIIS, Tsinghua
hehq13@mails.tsinghua.edu.cn, tianyiz16@mails.tsinghua.edu.cn

CP32

Communication Complexity of Discrete Fair Division

We initiate the study of the communication complexity of fair division with indivisible goods. We focus on some of the most well-studied fairness notions (envy-freeness, proportionality, and approximations thereof) and valuation classes (submodular, subadditive and unrestricted). Within these parameters, our results completely resolve whether the communication complexity of computing a fair allocation (or determining that none exist) is polynomial or exponential (in the number of goods), for every combination of fairness notion, valuation class, and number of players, for both deterministic and randomized protocols.

Benjamin Plaut
Stanford University
bplaut@stanford.edu

Tim Roughgarden
Stanford University
Computer Science Department
tim@cs.stanford.edu

CP32

Fully Polynomial-time Approximation Schemes for Fair Rent Division

We study the problem of fair rent division that entails splitting the rent and allocating the rooms of an apartment among roommates in a fair manner. In this setup, a distribution of the rent and an accompanying allocation is said to be fair if it is envy free, i.e., under the imposed rents, no agent has a strictly stronger preference for any other agent's room. The cardinal preferences of the agents are expressed via functions which specify the utilities of the agents for the rooms for every possible room rent/price. While envy-free solutions are guaranteed to exist under reasonably general utility functions, efficient algorithms for finding them were known only for quasilinear utilities. This work addresses this notable gap and develops approximation algorithms for fair rent division with minimal assumptions on the utility functions. Specifically, we show that if the agents have continuous, monotone decreasing, and piecewise-linear utilities, then the fair rent-division problem admits a fully polynomial-time approximation scheme. That is, we develop algorithms that find allocations and prices of the rooms such that for each agent a the utility of the room assigned to it is within a factor of $(1 + \epsilon)$ of the utility of the room most preferred by a; here, $\epsilon > 0$ is an approximation parameter. We complement the algorithmic results by proving that the fair rent division problem lies in the intersection of the complexity classes PPAD and PLS.

Eshwar Ram Arunachaleswaran
Indian Institute of Science
eshwarram.arunachaleswaran@gmail.com

Siddharth Barman
Department of Computer Sciences,
University of Wisconsin-Madison
barman@iisc.ac.in

Nidhi Rathi
Indian Institute of Science
nidhirathi@iisc.ac.in

CP32

An Optimal Truthful Mechanism for the Online Weighted Bipartite Matching Problem

In the weighted bipartite matching problem, the goal is to find a maximum-weight matching in a bipartite graph with nonnegative edge weights. We consider its online version where the first vertex set is known beforehand, but vertices of the second set appear one after another. Vertices of the first set are interpreted as items, and those of the second set as bidders. On arrival, each bidder vertex reveals the weights of all adjacent edges and the algorithm has to decide which of those to add to the matching. We introduce an optimal, e-competitive truthful mechanism under the

assumption that bidders arrive in random order (secretary model). It has been shown that the upper and lower bound of e for the original secretary problem extends to various other problems even with rich combinatorial structure, one of them being weighted bipartite matching. But truthful mechanisms so far fall short of reasonable competitive ratios once respective algorithms deviate from the original, simple threshold form. The best known mechanism for weighted bipartite matching offers only a ratio logarithmic in the number of online vertices. We close this gap, showing that truthfulness does not impose any additional bounds. The proof technique is new in this surrounding, and based on the observation of an independency inherent to the mechanism. The insights provided hereby are interesting in their own right and appear to offer promising tools for other problems, with or without truthfulness.

Rebecca Reiffenhaeuser
RWTH Aachen University
rebeccar@diag.uniroma1.it

CP32

Prophet Secretary Through Blind Strategies

In the classic prophet inequality, a problem in optimal stopping theory, samples from independent random variables arrive online. A gambler that knows the distributions, but cannot see the future, must decide at each point in time whether to stop and pick the current sample or to continue and lose that sample forever. The goal of the gambler is to maximize the expected value of what she picks and the performance measure is the worst case ratio between the expected value the gambler gets and what a prophet, that sees all the realizations in advance, gets. We study when the samples arrive in a uniformly random order, deriving a way of analyzing multi-threshold strategies that basically sets a nonincreasing sequence of thresholds to be applied at different times. The gambler will thus stop the first time a sample surpasses the corresponding threshold. We consider a class of robust strategies that we call blind quantile strategies. These constitute a clever generalization of single threshold strategies. Our main result shows that these strategies can achieve a constant of 0.669 in the prophet secretary problem, done by a sharp analysis of the underlying stopping time distribution for the gambler's strategy that is inspired by the theory of Schur-convex functions. We further prove that our family of blind strategies cannot lead to a constant better that 0.675. Finally we prove that no nonadaptive algorithm for the gambler can achieve a constant better than 0.732.

José Correa, Raimundo Saona
Universidad de Chile
correa@uchile.cl, raimundo.saona@ug.uchile.cl

Bruno Ziliotto
Université Paris Dauphine
CNRS
ziliotto@math.cnrs.fr

CP32

Pricing for Online Resource Allocation: Intervals and Paths

We present pricing mechanisms for several online resource allocation problems which obtain tight or nearly tight approximations to social welfare. In our settings, buyers arrive online and purchase bundles of items; buyers' values for the bundles are drawn from known distributions. This problem is closely related to the so-called prophet-inequality and its extensions in recent literature. Motivated by applications to cloud economics, we consider two kinds of buyer preferences. In the first, items correspond to different units of time at which a resource is available; the items are arranged in a total order and buyers desire intervals of items. The second corresponds to bandwidth allocation over a tree network; the items are edges in the network and buyers desire paths. For the interval preferences setting, we show that static, anonymous bundle pricings achieve a sublogarithmic competitive ratio, which is optimal (within constant factors) over the class of all online allocation algorithms, truthful or not. For the path preferences setting, we obtain a nearly-tight logarithmic competitive ratio. Both results exhibit an exponential improvement over item pricings for these settings. Our results extend to settings where the seller has multiple copies of each item, with the competitive ratio decreasing linearly with supply. Such a gradual tradeoff between supply and the competitive ratio for welfare was previously known only for the single item prophet inequality.

Shuchi Chawla, Benjamin Miller, Yifeng Teng
Department of Computer Sciences
University of Wisconsin-Madison
shuchi@cs.wisc.edu, bmiller@cs.wisc.edu, yifengt@cs.wisc.edu

CP33

Asymmetric Convex Intersection Testing

We consider *asymmetric convex intersection testing* (ACIT). Let $P \subset \mathbf{R}^d$ be a set of n points and $\mathcal{H}$ a set of n halfspaces in d dimensions. We denote by P the polytope obtained by taking the convex hull of P, and by $\mathcal{H}$ the polytope obtained by taking the intersection of the halfspaces in $\mathcal{H}$. Our goal is to decide whether the intersection of $\mathcal{H}$ and the convex hull of P are disjoint. Even though ACIT is a natural variant of classic LP-type problems that have been studied at length in the literature, and despite its applications in the analysis of high-dimensional data sets, it appears that the problem has not been studied before. We discuss how known approaches can be used to attack the ACIT problem, and we provide a very simple strategy that leads to a deterministic algorithm, linear on n and m, whose running time depends reasonably on the dimension d.

Luis Barba
ETH Zürich
luis.barba@inf.ethz.ch

Wolfgang J. Mulzer
Freie Universität Berlin
mulzer@inf.fu-berlin.de

CP33

Relaxed Voronoi: A Simple Framework for Terminal-Clustering Problems

We reprove three known algorithmic bounds for terminal-clustering problems, using a single framework that leads to simpler proofs. In this genre of problems, the input is a metric space (X, d) and a subset of terminals $K \subset X$, and the goal is to partition the points X such that each part, called a cluster, contains exactly one terminal (possibly with connectivity constraints) so as to minimize some goal. The bounds we reprove are for Steiner Point Removal on trees [Gupta, SODA 2001], for Metric 0-Extension in

bounded doubling dimension [Lee and Naor, unpublished 2003], and for Connected Metric 0-Extension [Englert et al., SICOMP 2014]. A natural approach is to cluster each point with its closest terminal, which would partition X into so-called Voronoi cells, but this approach can fail miserably due to its stringent cluster boundaries. A now-standard fix, which we call the Relaxed-Voronoi framework, is to use enlarged Voronoi cells, but to obtain disjoint clusters, the cells are computed greedily according to some order. This method, first proposed by Calinescu, Karloff and Rabani [SICOMP 2004], was employed successfully to provide state-of-the-art results for terminal-clustering problems on general metrics. However, for restricted families of metrics, only more complicated, ad-hoc algorithms are known. Our main contribution is to demonstrate that the Relaxed-Voronoi algorithm is applicable to restricted metrics, and leads to simple algorithms and analyses.

Arnold Filtser
Ben-Gurion University of the Negev
arnold273@gmail.com

Robert Krauthgamer, Ohad Trabelsi
Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il,
ohad.trabelsi@weizmann.ac.il

CP33

Approximating Optimal Transport with Linear Programs

In the regime of bounded transportation costs, additive approximations for the optimal transport problem are reduced (rather simply) to relative approximations for positive linear programs, resulting in faster additive approximation algorithms for optimal transport.

Kent Quanrud
University of Illinois at Urbana-Champaign
quanrud2@illinois.edu

CP33

On Primal-Dual Circle Representations

The Koebe-Andreev-Thurston Circle Packing Theorem states that every triangulated planar graph has a contact representation by circles. The theorem has been generalized in various ways. The most prominent generalization assures the existence of a primal-dual circle representation for every 3-connected planar graph. We present a simple and elegant elementary proof of this result.

Stefan Felsner
Technical University of Berlin
Institute for Mathematics
felsner@math.tu-berlin.de

Günter Rote
Freie Universität Berlin
Institut für Informatik
rote@inf.fu-berlin.de

CP33

LP Relaxation and Tree Packing for Minimum k-Cuts

Karger used spanning tree packings to derive a near linear-time randomized algorithm for the global minimum cut problem as well as a bound on the number of approximate minimum cuts. This is a different approach from his well-known random contraction algorithm. Thorup developed a fast deterministic algorithm for the minimum k-cut problem via greedy *recursive* tree packings. In this paper, we revisit the properties of an LP relaxation for k-cut proposed by Naor and Rabani, and analyzed by Chekuri, Guha and Naor. We show that the dual of the LP yields a tree packing, that when combined with an upper bound on the integrality gap for the LP, easily and transparently extends Karger's analysis for mincut to the k-cut problem. In addition to the simplicity of the algorithm and its analysis, this allows us to improve the running time of Thorup's algorithm by a factor of n. We also improve the bound on the number of α-approximate k-cuts. Second, we give a simple proof that the integrality gap of the LP is $2(1-1/n)$. Third, we show that an optimum solution to the LP relaxation, for all values of k, is fully determined by the principal sequence of partitions of the input graph. This allows us to relate the LP relaxation to the Lagrangean relaxation approach of Barahona and Ravi and Sinha; it also shows that the idealized recursive tree packing considered by Thorup gives an optimum dual solution to the LP.

Chandra Chekuri, Kent Quanrud
University of Illinois at Urbana-Champaign
chekuri@illinois.edu, quanrud2@illinois.edu

Chao Xu
Yahoo! Research NYC
chao.xu@oath.com

CP34

On the Structure of Unique Shortest Paths in Graphs

This paper develops a structural theory of unique shortest paths in real-weighted graphs. Our main goal is to characterize exactly which sets of node sequences, which we call path systems, can appear as unique shortest paths in a graph with arbitrary real edge weights. We say that such a path system is strongly metrizable. An easy fact implicit in the literature is that a strongly metrizable path system must be consistent, meaning that no two of its paths may intersect, split apart, and then intersect again. Our main result characterizes strong metrizability via forbidden intersection patterns along these lines. In other words, we describe a new family of forbidden patterns beyond consistency, and we prove that a path system is strongly metrizable if and only if it consistent and it avoids all of these new patterns. We offer separate (but closely related) characterizations in this way for the settings of directed, undirected, and directed acyclic graphs. Our characterizations are based on a new connection between shortest paths and certain boundary operators used in homology, which is used to prove several additional structural corollaries and seems to suggest new and possibly deep-rooted connections between shortest paths and topology.

Greg Bodwin
MIT
greg.bodwin@gmail.com

CP34

Exact Distance Oracles for Planar Graphs with Failing Vertices

We consider exact distance oracles for directed weighted planar graphs in the presence of failing vertices. Given a

source vertex u, a target vertex v and a set X of k failed vertices, such an oracle returns the length of a shortest u-to-v path that avoids all vertices in X. We propose oracles that can handle any number k of failures. More specifically, for a directed weighted planar graph with n vertices, any constant k, and for any $q \in [1, \sqrt{n}]$, we propose an oracle of size $\tilde{\mathcal{O}}(\frac{n^{k+3/2}}{q^{2k+1}})$ that answers queries in $\tilde{\mathcal{O}}(q)$ time. In particular, we show an $\tilde{\mathcal{O}}(n)$-size, $\tilde{\mathcal{O}}(\sqrt{n})$-query-time oracle for any constant k. This matches, up to polylogarithmic factors, the fastest failure-free distance oracles with nearly linear space. For single vertex failures ($k = 1$), our $\tilde{\mathcal{O}}(\frac{n^{5/2}}{q^3})$-size, $\tilde{\mathcal{O}}(q)$-query-time oracle improves over the previously best known tradeoff of Baswana et al. [SODA 2012] by polynomial factors for $q = \Omega(n^t)$, $t \in (1/4, 1/2]$. For multiple failures, no planarity exploiting results were previously known.

Panagiotis Charalampopoulos
Department of Informatics, Kings College London
panagiotis.charalampopoulos@kcl.ac.uk

Shay Mozes
IDC Herzliya, Israel
smozes@idc.ac.il

Benjamin Tebeka
Efi Arazi School of Computer Science, The Interdisciplinary
Center Herzliya
benitbk@gmail.com

CP34

Near Optimal Algorithms For The Single Source Replacement Paths Problem

The Single Source Replacement Paths (SSRP) problem is as follows; Given a graph $G = (V, E)$, a source vertex s and a shortest paths tree T_s rooted in s, output for every vertex $t \in V$ and for every edge e in T_s the length of the shortest path from s to t avoiding e. We present near optimal upper bounds, by providing $\tilde{O}(m\sqrt{n}+n^2)$ time randomized combinatorial algorithm (where n is the number of vertices, m is the number of edges and the $\tilde{O}$ notation suppresses polylogarithmic factors) for unweighted undirected graphs, and matching conditional lower bounds for the SSRP problem.

Sarel Cohen
Tel-Aviv University
sarelco@gmail.com

Shiri Chechik
Tel-Aviv University, Israel
shiri.chechik@gmail.com

CP34

The I/O Complexity of Toom-Cook Integer Multiplication

Nearly matching upper and lower bounds are derived for the I/O complexity of the Toom-Cook-k algorithm computing the products of two integers, each represented with n digits in a given base s, in a two-level storage hierarchy with M words of fast memory, with different digits stored in different memory words. An $IO_{\mathcal{A}_k}(n, M) = \Omega\left((n/M)^{\log_k 2k-1} M\right)$ lower bound on the I/O complexity is established, by a technique that combines an analysis of the size of the dominators of suitable sub-CDAGs of the Toom-Cook-k CDAG (Computational Directed Acyclic Graph) and the analysis of a function, which we call "*Partial Grigoriev's flow*', which captures the amount of information to be transferred between specific subsets of input and output variables, by any algorithm that solves the integer multiplication problem. The lower bound applies even if the recomputation of partial results is allowed. A careful implementation of the Toom-Cook-k algorithm, assuming that $M = \Omega\left(k^3 \log_s k\right)$, is also developed and analyzed, leading to an I/O complexity upper bound that is within a factor $O(k^2)$ of the corresponding lower bound, hence asymptotically optimal for fixed k. All bounds are actually derived in the more general case where the value of k is allowed to vary with the level of recursion. Extensions of the lower bound for a parallel model with P processors are also discussed.

Lorenzo De Stefani
Brown University
lorenzo_destefani@brown.edu

Gianfranco Bilardi
University of Padova
bilardi@dei.unipd.it

CP34

I/O-Efficient Algorithms for Topological Sort and Related Problems

This paper presents I/O-efficient algorithms for topologically sorting a directed acyclic graph and for the more general problem identifying and topologically sorting the strongly connected components of a directed graph $G = (V, E)$. Both algorithms are randomized and have I/O-costs $O(sort(E) \cdot \text{poly}(\lg V))$, with high probability, where $sort(E) = O(\frac{E}{B} \log_{M/B}(E/B))$ is the I/O cost of sorting an $|E|$-element array on a machine with size-B blocks and size-M cache/internal memory. These are the first algorithms for these problems that do not incur at least one I/O per vertex, and as such these are the first I/O-efficient algorithms for sparse graphs. By applying the technique of time-forward processing, these algorithms also imply I/O-efficient algorithms for most problems on directed acyclic graphs, such as shortest paths, as well as the single-source reachability problem on arbitrary directed graphs.

Nairen Cao, Jeremy Fineman, Katina Russell, Eugene Yang
Georgetown University
nc645@georgetown.edu, jfineman@cs.georgetown.edu, katina.russell@cs.georgetown.edu, eugene@ir.cs.georgetown.edu

CP35

Analyzing Boolean Functions on the Biased Hypercube Via Higher-Dimensional Agreement Tests

We propose a new paradigm for studying the structure of Boolean functions on the *biased* Boolean hypercube, i.e. when the measure is μ_p and p is potentially very small, e.g. as small as $O(1/n)$. Our paradigm is based on the following simple fact: the p-biased hypercube is expressible as a convex combination of many small-dimensional copies of the uniform hypercube. To uncover structure for μ_p, we invoke known structure theorems for $\mu_{1/2}$, obtaining a structured approximation for each copy separately. We then sew these approximations together using a novel

"agreement theorem". This strategy allows us to lift structure theorems from $\mu_{1/2}$ to μ_p. Our main application is a structure theorem for functions that are nearly low degree in the Fourier sense. The structure we uncover in the biased hypercube is *not at all* the same as for the uniform hypercube, despite using the structure theorem for the uniform hypercube as a black box. Rather, new phenomena emerge: whereas nearly low degree functions on the uniform hypercube are close to juntas, when p becomes small, non-juntas arise as well. A key component of our proof is a new local-to-global agreement theorem for higher dimensions, which extends the work of Dinur and Steurer [*Proc. 29th CCC*, 2014]. Whereas their result sews together vectors, our agreement theorem sews together labeled graphs and hypergraphs.

Yuval Filmus
Technion
yuvalfi@cs.technion.ac.il

Irit Dinur
Weizmann Institute
irit.dinur@weizmann.ac.il

Prahladh Harsha
Tata institute
prahladh@tifr.res.in

CP35

Maximally Recoverable LRCs: A Field Size Lower Bound and Constructions for Few Heavy Parities

The explosion in the volumes of data being stored online has resulted in distributed storage systems transitioning to erasure coding based schemes. Local Reconstruction Codes (LRCs) have emerged as the codes of choice for these applications. These codes can correct a small number of erasures by accessing only a small number of remaining coordinates. An (n, r, h, a, q)-LRC is a linear code over $\mathbb{F}_q$ of length n, whose codeword symbols are partitioned into $g = n/r$ local groups each of size r. It has h global parity checks and each local group has a local parity checks. Such an LRC is Maximally Recoverable (MR), if it corrects all erasure patterns which are information-theoretically correctable given the structure of local and global parity checks. We show that when a and h are constant and r may grow with n, for every MR LRC with $g = n/r$ local groups,

$$q \geq \Omega_{a,h}\left(n \cdot r^{\alpha}\right) \text{ where } \alpha = \frac{\min\{a, h - 2\lceil h/g \rceil\}}{\lceil h/g \rceil}.$$

No superlinear (in n) lower bounds were known prior to this work for any setting of parameters. MR LRCs deployed in practice have a small number of global parities, typically $h = 2, 3$. We complement our lower bounds by giving constructions with small field size for $h \leq 3$. For $h = 2$, we give a linear field size construction. We also show a surprising application of elliptic curves and arithmetic progression free sets in the construction of MR LRCs.

Sivakanth Gopi
Microsoft Research
sigopi@microsoft.com

Venkatesan Guruswami
Carnegie Mellon University
venkatg@cs.cmu.edu

Sergey Yekhanin
Microsoft
yekhanin@microsoft.com

CP35

List Decoding with Double Samplers

Samplers are bipartite graphs with nice pseudo-random properties. A classical construction of ABNNR uses sampler graphs for amplifying the distance of an error correcting code. This construction has an algorithm for unique decoding, but no efficient list decoding algorithm is known. In this talk, I will introduce "double samplers", which extend sampler graphs to three layers, and show a polynomial time algorithm which list decodes the ABNNR code construction on a double sampler. Double samplers can be constructed from high dimensional expanders, and currently no other method is known.

Inbal R. Livni Navon, Irit Dinur
Weizmann Institute
inbal.livni@weizmann.ac.il, irit.dinur@weizmann.ac.il

Prahladh Harsha
Tata institute
prahladh@tifr.res.in

Tali Kaufman
MIT
kaufmant@mit.edu

Amnon Ta-Shma
Tel-Aviv University
amnon@tau.ac.il

CP35

Binary Robust Positioning Patterns with Low Redundancy and Efficient Locating Algorithms

A robust positioning pattern is a large array that allows a mobile device to locate its position by reading a possibly corrupted small window around it. This paper provides constructions of binary positioning patterns, equipped with efficient locating algorithms, that are robust to a constant number of errors and have redundancy within a constant factor of optimality. Our construction of binary robust positioning sequences has the least known redundancy amongst those explicit constructions with efficient locating algorithms. On the other hand, for binary robust positioning arrays, our construction is the first explicit construction whose redundancy is within a constant factor of optimality. The locating algorithms accompanying our constructions run in time cubic in sequence length or array dimensions.

Yeow Meng Chee, Duc Tu Dao
School of Physical and Mathematical Sciences
Nanyang Technological University
ymchee@ntu.edu.sg, daoductu001@ntu.edu.sg

Han Mao Kiah
NTU, Singapore
hmkiah@ntu.edu.sg

San Ling, Hengjia Wei
School of Physical and Mathematical Sciences
Nanyang Technological University

lingsan@ntu.edu.sg, hjwei@ntu.edu.sg

CP35

Synchronization Strings: Highly Efficient Deterministic Constructions over Small Alphabets

Synchronization strings are recently introduced by Haeupler and Shahrasbi (STOC17) in the study of codes for correcting insertion and deletion errors. A synchronization string is an encoding of the indices of the symbols in a string, and together with an appropriate decoding algorithm it can transform insertion and deletion errors into standard symbol erasures and corruptions. This reduces the problem of constructing insdel codes to the problem of constructing standard error correcting codes. Besides this, synchronization strings are also useful in other applications such as synchronization sequences and interactive coding schemes. For all such applications, synchronization strings are desired to be over alphabets that are as small as possible. Haeupler and Shahrasbi showed that for any parameter $\varepsilon > 0$, synchronization strings of arbitrary length exist over an alphabet whose size depends only on ε. Specifically, they obtained an alphabet size of $O(\varepsilon^{-4})$, which left an open question on where the minimal size of such alphabets lies between $\Omega(\varepsilon^{-1})$ and $O(\varepsilon^{-4})$. In this work, we partially bridge this gap by providing an improved lower bound of $\Omega\left(\varepsilon^{-3/2}\right)$, and an improved upper bound of $O\left(\varepsilon^{-2}\right)$. We also provide fast explicit constructions of synchronization strings over small alphabets.

Kuan Cheng
Johns Hopkins University
kcheng17@jhu.edu

Bernhard Haeupler
Carnegie Mellon U.
haeupler@cs.cmu.edu

Xin Li
Johns Hopkins University
lixints@cs.jhu.edu

Amirbehshad Shahrasbi
Carnegie Mellon University
shahrasbi@cs.cmu.edu

Ke Wu
Johns Hopkins University
kwu35@jhu.edu

CP36

Improved Bounds for Randomly Sampling Colorings Via Linear Programming

A well-known conjecture in computer science and statistical physics is that Glauber dynamics on the set of k-colorings of a graph G on n vertices with maximum degree Δ is rapidly mixing for $k \geq \Delta + 2$. In FOCS 1999, Vigoda showed that the flip dynamics (and therefore also Glauber dynamics) is rapidly mixing for any $k > \frac{11}{6}\Delta$. It turns out that there is a natural barrier at $\frac{11}{6}$, below which there is no one-step coupling that is contractive with respect to the Hamming metric, even for the flip dynamics. We use linear programming and duality arguments to fully characterize the obstructions to going beyond $\frac{11}{6}$. These extremal configurations turn out to be quite brittle, and in this paper we use this to give two proofs that the Glauber dynamics is rapidly mixing for any $k \geq (\frac{11}{6} - \epsilon_0)\Delta$ for some absolute constant $\epsilon_0 > 0$. This is the first improvement to Vigoda's result that holds for general graphs. Our first approach analyzes a variable-length coupling in which these configurations break apart with high probability before the coupling terminates, and our other approach analyzes a one-step path coupling with a new metric that counts the extremal configurations. Additionally, our results extend to list coloring, a widely studied generalization of coloring, where the previously best known results required $k > 2\Delta$.

Sitan Chen
MIT
sitanc@mit.edu

Michelle Delcourt
University of Waterloo
mdelcour@uwaterloo.ca

Ankur Moitra
MIT
moitra@mit.edu

Guillem Perarnau
University of Birmingham
g.perarnau@bham.ac.uk

Luke Postle
University of Waterloo
lpostle@math.uwaterloo.ca

CP36

The Complexity of Approximately Counting Retractions

Let G be a graph that contains an induced subgraph H. A *retraction* from G to H is a homomorphism from G to H that is the identity function on H. Retractions are very well-studied: Given H, the complexity of deciding whether there is a retraction from an input graph G to H is completely classified, in the sense that it is known for which H this problem is tractable (assuming P $\neq$ NP). Similarly, the complexity of (exactly) counting retractions from G to H is classified (assuming FP $\neq$ #P). However, almost nothing is known about approximately counting retractions. Our first contribution is to give a complete trichotomy for approximately counting retractions to trees. Our second contribution is to locate the retraction counting problem in the complexity landscape of related approximate counting problems. Interestingly, our results are in contrast to the situation in the exact counting context. We show that the problem of approximately counting retractions is separated both from the problem of approximately counting homomorphisms and from the problem of approximately counting list homomorphisms — whereas for exact counting all three of these problems are interreducible. We also show that the number of retractions is at least as hard to approximate as both the number of surjective homomorphisms and the number of compactions. In contrast, exactly counting compactions is the hardest of these problems.

Jacob Focke, Leslie Ann Goldberg, Stanislav Zivny
University of Oxford
jacob.focke@cs.ox.ac.uk, leslie.goldberg@cs.ox.ac.uk,

standa.zivny@cs.ox.ac.uk

CP36

Algorithms for #BIS-Hard Problems on Expander Graphs

We give an FPTAS and an efficient sampling algorithm for the high-fugacity hard-core model on bounded-degree bipartite expander graphs and the low-temperature ferromagnetic Potts model on bounded-degree expander graphs. The results apply, for example, to random (bipartite) Δ-regular graphs, for which no efficient algorithms were known for these problems (with the exception of the Ising model) in the non-uniqueness regime of the infinite Δ-regular tree.

Matthew Jenssen, Peter Keevash
University of Oxford
jenssen@maths.ox.ac.uk, keevash@math.ox.ac.uk

Will Perkins
University of Birmingham
math@willperkins.org

CP36

Approximability of the Six-Vertex Model

Six-vertex models originate in statistical mechanics, as a family of vertex models for crystal lattices with hydrogen bonds. In the language of graph theory and theory of computing, it is a sum-of-product computation. On a 4-regular graph, we compute the partition function which is a weighted sum of Eulerian orientations, where at every vertex the orientation must be two-in-two-out (called the *ice rule*). There are thousands of papers on the six-vertex model, making it one of the three most studied models in statistical physics, together with ferromagnetic Ising and monomer-dimer models. In this paper, we take the first step toward a classification of the approximation complexity of the six-vertex model. Our complexity results conform to the phase transition phenomenon from physics. We show that the approximation complexity of the six-vertex model behaves dramatically differently on the two sides separated by the phase transition threshold. Furthermore, we present structural properties of the six-vertex model on planar graphs for parameter settings that have known relations to the Tutte polynomial $T(G; x, y)$. In this talk, I will outline our main results and techniques in this work.

Jin-Yi Cai
University of Wisconsin, Madison
jyc@cs.wisc.edu

Tianyu Liu
University of WisconsinMadison
tl@cs.wisc.edu

Pinyan Lu
Shanghai University of Finance and Economics
lu.pinyan@mail.shufe.edu.cn

CP36

Zeros of Holant Problems: Locations and Algorithms

We present fully polynomial-time (deterministic or randomised) approximation schemes for Holant problems, defined by a non-negative constraint function satisfying a generalised second order recurrence modulo a couple of exceptional cases. As a consequence, any non-negative Holant problem on cubic graphs has an efficient approximation algorithm unless the problem is equivalent to approximately counting perfect matchings, a central open problem in the area. This is in sharp contrast to the computational phase transition shown by 2-state spin systems on cubic graphs. Our main technique is the recently established connection between zeros of graph polynomials and approximate counting. We also use the winding technique to deduce the second result on cubic graphs.

Heng Guo
University of Edinburgh
hguo@inf.ed.ac.uk

Chao Liao
Shanghai Jiao Tong University
chao.liao.95@gmail.com

Pinyan Lu
Shanghai University of Finance and Economics
lu.pinyan@mail.shufe.edu.cn

Chihao Zhang
Shanghai Jiaotong University
chihao.zhang@gmail.com

CP37

Towards a Unified Theory of Sparsification for Matching Problems

We present a construction of a "matching sparsifier', that is, a sparse subgraph of the given graph that preserves large matchings approximately and is robust to modifications of the graph. We use this matching sparsifier to obtain several new algorithmic results for the maximum matching problem:

- An almost (3/2)-approximation one-way communication protocol for the maximum matching problem, simplifying the (3/2)-approximation protocol of Goel, Kapralov, and Khanna (SODA 2012) and extending it from bipartite graphs to general graphs.
- An almost (3/2)-approximation algorithm for the stochastic matching problem, improving upon and simplifying the previous 1.999-approximation algorithm of Assadi, Khanna, and Li (EC 2017).
- An almost (3/2)-approximation algorithm for the fault-tolerant matching problem, implying the first non-trivial algorithm for this problem.

Our matching sparsifier is obtained by proving new properties of the edge-degree constrained subgraph (EDCS) of Bernstein and Stein (ICALP 2015; SODA 2016)—designed in the context of matchings in dynamic graphs—that identifies EDCS as an excellent choice for a matching sparsifier. This leads to surprisingly simple and non-technical proofs of the above results in a unified way. Along the way, we also provide a simpler proof of the fact that an EDCS is guaranteed to contain a large matching, which may be of independent interest.

Sepehr Assadi
University of Pennsylvania
sassadi@cis.upenn.edu

Aaron Bernstein
Technical University of Berlin, Germany

bernstei@gmail.com

CP37

A New Application of Orthogonal Range Searching for Computing Giant Graph Diameters

A well-known problem for which it is difficult to improve the textbook algorithm is computing the graph diameter. We present two versions of a simple algorithm (one being Monte Carlo and the other deterministic) that for every fixed h and graph G, either correctly concludes that $diam(G) < hn$ or outputs $diam(G)$, in time $\mathcal{O}(m + n^{1+o(1)})$. The algorithm combines a simple randomized strategy for this problem (Damaschke, *IWOCA'16*) with a popular framework for computing graph distances that is based on range trees (Cabello and Knauer, *Computational Geometry'09*). We also prove that under the Strong Exponential Time Hypothesis (SETH), we cannot compute the diameter of a given n-vertex graph in truly subquadratic time, even if the diameter is an $\Theta(n/\log n)$.

Guillaume Ducoffe
ICI National Institute for Research and Development in Informatics
guillaume.ducoffe@ici.ro

CP37

A Note on Max K-Vertex Cover: Faster FPT-AS, Smaller Approximate Kernel and Improved Approximation

In Maximum k-Vertex Cover (Max k-VC), the input is an edge-weighted graph G and an integer k, and the goal is to find a subset S of k vertices that maximizes the total weight of edges covered by S. (An edge is covered by S iff at least one of its endpoints lies in S.) We present a simple FPT approximation scheme (FPT-AS) that runs in $(1/\epsilon)^{O(k)}poly(n)$ time for the problem, which improves upon Gupta et al.'s $(k/\epsilon)^{O(k)}poly(n)$-time FPT-AS [SODA'18, FOCS'18]. Our algorithm naturally yields an approximate kernelization scheme of $O(k/\epsilon)$ vertices; previously, an $O(k^5/\epsilon^2)$-vertex approximate kernel is only known for unweighted Max k-VC [Lokshtanov et al., STOC'17]. Moreover, using our approximate kernelization as a preprocessing step, we can directly apply Raghavendra and Tan's algorithm for 2SAT with cardinality constraint [SODA'12] to give an 0.92-approximation algorithm for Max k-VC in polynomial time. This improves upon Feige and Langberg's algorithm [J. Algorithms'01] which yields $(0.75 + \delta)$-approximation for some $\delta > 0$. We also consider the minimization version (called Min k-VC), where the goal is to minimize the total weight of edges covered by S. We provide an FPT-AS for Min k-VC with similar running time of $(1/\epsilon)^{O(k)}poly(n)$. On the other hand, we show that there is unlikely a polynomial size approximate kernelization for Min k-VC for any factor less than two.

Pasin Manurangsi
UC Berkeley
pasin@berkeley.edu

CP37

Simple Greedy 2-Approximation Algorithm for the Maximum Genus of a Graph

The maximum genus $\gamma_M(G)$ of a graph G is the largest genus of an orientable surface into which G has a cellular embedding. Combinatorially, it coincides with the maximum number of disjoint pairs of adjacent edges of G whose removal results in a connected spanning subgraph of G. In this paper we describe a greedy 2-approximation algorithm for maximum genus by proving that removing pairs of adjacent edges from G arbitrarily while retaining connectedness leads to at least $\gamma_M(G)/2$ pairs of edges removed. As a consequence of our approach we also obtain a 2-approximate counterpart of Xuong's combinatorial characterisation of maximum genus.

Michal Kotrbcik
Comenius University, Department of Computer Science
Mlynska Dolina, 842 48 Bratislava
kotrbcik@dcs.fmph.uniba.sk

Martin Skoviera
Comenius University, Bratislava, Slovakia
skoviera@dcs.fmph.uniba.sk

CP37

Simplified and Space-Optimal Semi-Streaming $(2+\epsilon)$-Approximate Matching

In a recent breakthrough, Paz and Schwartzman (SODA'17) presented a single-pass $(2 + \epsilon)$-approximation algorithm for the maximum weight matching problem in the semi-streaming model. Their algorithm uses $O(n \log^2 n)$ bits of space, for any constant $\epsilon > 0$. We present a simplified and more intuitive primal-dual analysis, for essentially the same algorithm, which also improves the space complexity to the optimal bound of $O(n \log n)$ bits — this is optimal as the output matching requires $\Omega(n \log n)$ bits.

Mohsen Ghaffari
ETH Zurich
ghaffari@inf.ethz.ch

David Wajc
Carnegie Mellon University
dwajc@cs.cmu.edu

CP38

Hierarchical Clustering Better than Average-Linkage

Hierarchical Clustering (HC) is a widely studied problem in exploratory data analysis, usually tackled by simple agglomerative procedures like average-linkage, single-linkage or complete-linkage. In this paper we focus on two objectives, introduced recently to give insight into the performance of average-linkage clustering: a similarity based HC objective proposed by [Moseley and Wang, 2017] and a dissimilarity based HC objective proposed by [Cohen-Addad et al., 2018]. In both cases, we present tight counterexamples showing that average-linkage cannot obtain better than 1/3 and 2/3 approximations respectively (in the worst-case), settling an open question raised in [Moseley and Wang, 2017]. This matches the approximation ratio of a random solution, raising a natural question: can we beat average-linkage for these objectives? We answer this in the affirmative, giving two new algorithms based on semidefinite programming with provably better guarantees.

Moses Charikar, Vaggos Chatziafratis, Rad Niazadeh
Stanford University, California
moses@cs.stanford.edu, vaggos@cs.stanford.edu,

rad@cs.stanford.edu

CP38

Universal Trees Grow Inside Separating Automata: Quasi-Polynomial Lower Bounds for Parity Games

Several distinct techniques have been proposed to design quasi-polynomial algorithms for solving parity games since the breakthrough result of Calude, Jain, Khoussainov, Li, and Stephan (2017): play summaries, progress measures and register games. We argue that all those techniques can be viewed as instances of the separation approach to solving parity games, a key technical component of which is constructing (explicitly or implicitly) an automaton that separates languages of words encoding plays that are (decisively) won by either of the two players. Our main technical result is a quasi-polynomial lower bound on the size of such separating automata that nearly matches the current best upper bounds. This forms a barrier that all existing approaches must overcome in the ongoing quest for a polynomial-time algorithm for solving parity games. The key and fundamental concept that we introduce and study is a universal ordered tree. The technical highlights are a quasi-polynomial lower bound on the size of universal ordered trees and a proof that every separating safety automaton has a universal tree hidden in its state space.

Wojciech Czerwinski
University of Warsaw
wczerwin@mimuw.edu.pl

Laure Daviaud
University of Warwick
l.daviaud@warwick.ac.uk

Nathanael Fijalkow
CNRS, University of Bordeaux
nathanael.fijalkow@labri.fr

Marcin Jurdzinski, Ranko Lazic
University of Warwick
marcin.jurdzinski@warwick.ac.uk,
R.S.Lazic@warwick.ac.uk

Pawel Parys
University of Warsaw
parys@mimuw.edu.pl

CP38

On Facility Location with General Lower Bounds

In this paper, we give the first constant approximation algorithm for the lower bounded facility location (LBFL) problem with general lower bounds. Prior to our work, such algorithms were only known for the special case where all facilities have the same lower bound: Svitkina [Svi10] gave a 448-approximation for the special case, and subsequently Ahmadian and Swamy [AS13] improved the approximation factor to 82.6. As in [Svi10] and [AS13], our algorithm for LBFL with general lower bounds works by reducing the problem to the capacitated facility location (CFL) problem. To handle the challenges raised by the general lower bounds, it involves more reduction steps. One main complication is that after aggregating the clients and facilities at a few locations, each of these locations may contain many facilities with different opening costs and lower bounds. To address this issue, we introduce and reduce the LBFL problem to two intermediate problems called the LBFL with penalty (LBFL-P) and the transportation with configurable supplies and demands (TCSD) problems, which in turn can be reduced to the CFL problem.

Shi Li
Toyota Technological Institute at Chicago
shil@buffalo.edu

CP38

Exponential Lower Bounds on Spectrahedral Representations of Hyperbolicity Cones

In an effort to better understand the relationship between semidefinite programming and hyperbolic programming, we study dimension implications of certain representations. The Generalized Lax Conjecture asks whether every hyperbolicity cone is a section of a semidefinite cone of sufficiently high dimension. We prove that the space of hyperbolicity cones of hyperbolic polynomials contains pairwise distant cones in the Hausdorff metric, and therefore provide lower bounds on the size of representations of these polynomials (even allowing a small approximation error). The cones are perturbations of the hyperbolicity cones of elementary symmetric polynomials.

Nick Ryder
UC Berkeley
nick.ryder@berkeley.edu

Nikhil Srivastava
University of California Berkeley
nikhil@math.berkeley.edu

Prasad Raghavendra
U C Berkeley
prasad@cs.berkeley.edu

Benjamin Weitz
Pintrest
bsweitz123@gmail.com

CP38

The Threshold for SDP-Refutation of Random Regular NAE-3SAT

Unlike its cousin 3SAT, the NAE-3SAT (not-all-equal-3SAT) problem has the property that spectral/SDP algorithms can efficiently refute random instances when the constraint density is a large constant (with high probability). But do these methods work immediately above the "satisfiability threshold", or is there still a range of constraint densities for which random NAE-3SAT instances are unsatisfiable but hard to refute? We show that the latter situation prevails, at least in the context of random regular instances and SDP-based refutation. More precisely, whereas a random d-regular instance of NAE-3SAT is easily shown to be unsatisfiable (whp) once $d \geq 8$, we establish the following sharp threshold result regarding efficient refutation: If $d < 13.5$ then the basic SDP, even augmented with triangle inequalities, fails to refute satisfiability (whp), if $d > 13.5$ then even the most basic spectral algorithm refutes satisfiability (whp).

Tselil Schramm
UC Berkeley
tschramm@mit.edu

Yash Deshpande

MIT
yash@mit.edu

Andrea Montanari
Stanford University
montanari@stanford.edu

Ryan O'Donnell
Carnegie Mellon University, USA
odonnell@cs.cmu.edu

Subhabrata Sen
MIT
ssen90@mit.edu

CP39

Theorems of Carathodory, Helly, and Tverberg Without Dimension

Motivated by Barman, we initiate a systematic study of the 'no-dimensional' analogues of some basic theorems in combinatorial and convex geometry, including the colorful Carathéodory's theorem, Tverberg's theorem, Helly's theorem as well as their fractional and colorful extensions. We show that, besides optimally improving some of Barman's theorems, they have several algorithmic and combinatorial consequences.

Karim Adiprasito
Einstein Institute for Mathematics
Hebrew University of Jerusalem, Israel
adiprasito@math.huji.ac.il

Imre Bárány
Alfréd Rényi Institute of Mathematics, Budapest
University College London, UK
barany.imre@renyi.mta.hu

Nabil H. Mustafa
Laboratoire d'Informatique Gaspard-Monge
ESIEE Paris, France
mustafan@esiee.fr

Monika Csikos
Universite Paris-Est, Laboratoire d'Informatique
Gaspard-Mon
monika.csikos@esiee.fr

CP39

Minimizing Interference Potential Among Moving Entities

We consider the problem of monitoring the interference among a collection of entities moving with bounded speed in $\Re^d$. Uncertainty in entity locations due to unmonitored and unpredictable motion gives rise to a space of possible entity configurations at each time, with possibly very different interference properties. We define measures of the interference potential of such spaces to describe the interference that might actually occur. We study how limited monitoring rate impacts interference potential by studying a clairvoyant scheme (one that knows the trajectories of all entities) subject to the same monitoring restriction. This forms a benchmark for the analysis of uninformed schemes. We describe and analyse an adaptive monitoring scheme for minimizing interference potential over time that is competitive (to within a constant factor) with any other scheme (in particular, a clairvoyant scheme) over modest sized time intervals. Imagine mobile transmission sources with broadcast ranges that interfere if they transmit on the same channel and their ranges intersect. Location uncertainty effectively expands their broadcast range to a potential range. The chromatic number of the intersection graph of these potential ranges gives the minimum number of channels required to avoid interference. For fixed monitoring rate, our scheme provides an adaptive channel assignment that uses a number of channels competitive with any other scheme.

Daniel Busto
Ericsson
Montreal, Canada
danielsbusto@gmail.com

William S. Evans
University of British Columbia
will@cs.ubc.ca

David Kirkpatrick
University of British Columbia, Canada
kirk@cs.ubc.ca

CP39

On the Spanning and Routing Ratio of Theta-Four

We present a routing algorithm for the Θ_4-graph that computes a path between any two vertices s and t having length at most 17 times the Euclidean distance between s and t. At each step the algorithm only uses knowledge of the location of the current vertex, its (at most four) outgoing edges, the destination vertex, and one additional bit of information in order to determine the next edge to follow. This provides the first known online, local, competitive routing algorithm with constant routing ratio for the Θ_4-graph, as well as improving the best known upper bound on the spanning ratio from 237 to 17.

Darryl R. Hill, Prosenjit Bose
Carleton University
darryl.r.hill@gmail.com, jit@scs.carleton.ca

Jean-Lou De Carufel
University of Ottawa
jeanloudecarufel@gmail.com

Michiel Smid
Carleton University
michiel.smid@gmail.com

CP39

Greedy Spanners Are Optimal in Doubling Metrics

We show that the greedy spanner algorithm constructs a $(1+\epsilon)$-spanner of weight $\epsilon^{-O(d)}w(MST)$ for a point set in metrics of doubling dimension d, resolving an open problem posed by Gottlieb. Our result generalizes the result by Narasimhan and Smid who showed that a point set in d-dimension Euclidean space has a $(1+\epsilon)$-spanner of weight at most $\epsilon^{-O(d)}w(MST)$. Our proof only uses the packing property of doubling metrics and greatly simplifies the proof of the same result in Euclidean space.

Glencora Borradaile, Hung Le
Oregon State University
glencora@eecs.oregonstate.edu,
lehu@onid.oregonstate.edu

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen
koolooz@di.ku.dk

CP39

Viewing the Rings of a Tree: Minimum Distortion Embeddings into Trees

We describe a $(1+\varepsilon)$ approximation algorithm for finding the minimum distortion embedding of an n-point metric space, (X, X_d), into a tree with vertex set X. The running time is $n^2 \cdot (\Delta/\varepsilon)^{(O(\delta_{opt}/\varepsilon))^{2\lambda+1}}$ parameterized with respect to the spread of X, denoted by Δ, the minimum possible distortion for embedding X into any tree, denoted by δ_{opt}, and the doubling dimension of X, denoted by λ. Hence we obtain a PTAS, provided δ_{opt} is constant and X is a finite doubling metric space with polynomial spread, for example, a point set with polynomial spread in constant dimensional Euclidean space. Our algorithm implies a constant factor approximation with the same running time when Steiner vertices are allowed. Moreover, we describe a $(1+\varepsilon)$ approximation algorithm with a similar running time for finding a tree spanner of (X, X_d) that minimizes the maximum stretch. Finally, we generalize our tree spanner algorithm to a $(1+\varepsilon)$ approximation algorithm for computing a minimum stretch tree spanner of a weighted graph, where the running time is also parameterized with respect to maximum degree. In particular, we obtain a PTAS for computing minimum stretch tree spanners of weighted graphs, with polynomially bounded spread, constant doubling dimension, and constant maximum degree, when a tree spanner with constant stretch exists.

Benjamin A. Raichel
University of Texas at Dallas
benjamin.raichel@utdallas.edu

Amir Nayyeri
School of Electrical Eng and Computer Science
Oregon State University
nayyeria@eecs.oregonstate.edu

CP40

Towards Instance-optimal Private Query Release

We study efficient mechanisms for the query release problem in differential privacy: given a workload of m statistical queries, output approximate answers to the queries while satisfying the constraints of differential privacy. In particular, we are interested in mechanisms that optimally adapt to the given workload. Building on the projection mechanism of Nikolov, Talwar, and Zhang, and using the ideas behind Dudley's chaining inequality, we propose new efficient algorithms for the query release problem, and prove that they achieve optimal sample complexity for the given workload (up to constant factors, in certain parameter regimes) with respect to the class of mechanisms that satisfy concentrated differential privacy. We also give variants of our algorithms that satisfy local differential privacy, and prove that they also achieve optimal sample complexity among all local sequentially interactive private mechanisms.

Aleksandar Nikolov
University of Toronto
anikolov@cs.toronto.edu

Jaroslaw Blasiok
Harvard University
jblasiok@g.harvard.edu

Mark Bun
Princeton University
mbun@cs.princeton.edu

Thomas Steinke
IBM Research
opt-cdp@thomas-steinke.net

CP40

Lower Bounds for Oblivious Data Structures

An oblivious data structure is a data structure where the memory access patterns reveals no information about the operations performed on it. Such data structures were introduced by Wang *et al.* [ACM SIGSAC'14] and are intended for situations where one wishes to store the data structure at an untrusted server. One way to obtain an oblivious data structure is simply to run a classic data structure on an oblivious RAM (ORAM). With the current best ORAM implementations, this results in an overhead of $\omega(\lg n)$ for the most natural setting of parameters. Moreover, a recent lower bound for ORAMs by Larsen and Nielsen [CRYPTO'18] show that they always incur an overhead of at least $\Omega(\lg n)$ if used in a black box manner. To circumvent the current $\omega(\lg n)$ overhead, researchers have instead studied classic data structure problems more directly and have obtained efficient solutions for many such problems such as stacks, queues, deques, priority queues and search trees. However, none of these data structures process operations faster than $\Theta(\lg n)$, leaving open the question of whether even faster solutions exist. In this paper, we rule out this possibility by proving $\Omega(\lg n)$ lower bounds for oblivious stacks, queues, deques, priority queues and search trees.

Riko Jacob
IT University Copenhagen
rikj@itu.dk

Kasper G. Larsen
MADALGO, Aarhus University
larsen@cs.au.dk

Jesper B. Nielsen
Aarhus University
jbn@cs.au.dk

CP40

Foundations of Differentially Oblivious Algorithms

It is well-known that a program's memory access pattern can leak information about its input. To thwart such leakage, most existing works adopt the technique of oblivious RAM (ORAM) simulation. Such an obliviousness notion has stimulated much debate. Although ORAM techniques have significantly improved over the past few years, the concrete overheads are arguably still undesirable for real-world systems — part of this overhead is in fact inherent due to a well-known logarithmic ORAM lower bound by Goldreich and Ostrovsky. Inspired by the elegant notion of differential privacy, we initiate the study of a new notion of access pattern privacy, which we call "(ϵ, δ)-differential obliviousness'. We separate the notion of (ϵ, δ)-differential obliviousness from classical oblivious-

ness by considering several fundamental algorithmic abstractions including sorting small-length keys, merging two sorted lists, and range query data structures. We show that by adopting differential obliviousness with reasonable choices of ϵ and δ, not only can one circumvent several impossibilities pertaining to full obliviousness, one can also, in several cases, obtain meaningful privacy with little overhead relative to the non-private baselines.

T-H. Hubert Chan
The University of Hong Kong
hubert@cs.hku.hk

Kai-Min Chung
Academia Sinica
kmchung@iis.sinica.edu.tw

Bruce M. Maggs
Duke University and Akamai Technologies
bmm@cs.duke.edu

Elaine Shi
Cornell University
runting@gmail.com

CP40

Amplification by Shuffling: From Local to Central Differential Privacy via Anonymity

Sensitive statistics are often collected across sets of users, with repeated collection of reports done over time. For example, trends in users' private preferences or software usage may be monitored via such reports. Building on recent work, we study the collection of such statistics in the local differential privacy (LDP) model, and describe an algorithm whose privacy cost is polylogarithmic in the number of changes to a user's value. This algorithm is of particular practical benefit in applications where each user's value may change only a small number of times. More fundamentally—by building on anonymity of the users' reports—we also demonstrate how the privacy cost of our LDP algorithm can actually be much lower when viewed in the central model of differential privacy, We show, via a new and general technique, that any permutation-invariant algorithm satisfying ε-local differential privacy will satisfy $O(\varepsilon\log(1/\delta)/\sqrt{n}, \delta)$-central differential privacy. By this, we clarify how the high noise and $\sqrt{n}$ overhead of LDP protocols is a consequence of them being significantly more private in the central model. As a final, practical corollary, our results also imply that industrial deployments of LDP mechanism may have much lower privacy cost than their advertised ε would indicate—at least if reports are anonymized.

Ulfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar
Google
ulfar@google.com, vitalyfm@google.com, mironov@google.com, pseudorandom@google.com, kunal@google.com

Abhradeep Thakurta
UC Santa Cruz and Google
aguhatha@ucsc.edu

CP40

Can we Overcome the $n \log n$ Barrier for Oblivious Sorting?

It is well-known that non-comparison-based techniques can allow us to sort n elements in $o(n \log n)$ time on a Random-Access Machine (RAM). On the other hand, it is a long-standing open question whether (non-comparison-based) circuits can sort n elements from the domain $[1..2^k]$ with $o(kn \log n)$ boolean gates. We consider weakened forms of this question: first, we consider a restricted class of sorting where the number of distinct keys is much smaller than the input length; and second, we explore Oblivious RAMs and probabilistic circuit families, i.e., computational models that are somewhat more powerful than circuits but much weaker than RAM. We show that Oblivious RAMs and probabilistic circuit families can sort $o(\log n)$-bit keys in $o(n \log n)$ time or $o(kn \log n)$ circuit complexity. Our algorithms work in the balls-and-bins model, i.e., not only can they sort an array of numerical keys — if each key additionally carries an opaque ball, our algorithms can also move the balls into the correct order. We further show that in such a balls-and-bins model, it is impossible to sort $\Omega(\log n)$-bit keys in $o(n \log n)$ time, and thus the $o(\log n)$-bit-key assumption is necessary for overcoming the $n \log n$ barrier. Finally, after optimizing the IO efficiency, we show that even the 1-bit special case can solve open questions: our oblivious algorithms solve tight compaction and selection with optimal IO efficiency for the first time.

Wei-Kai Lin, Elaine Shi
Cornell University
wklin@cs.cornell.edu, elaine@cs.cornell.edu

Tiancheng Xie
Shanghai Jiao Tong University
niconiconi@sjtu.edu.cn

CP41

A Simple Near-Linear Pseudopolynomial Time Randomized Algorithm for Subset Sum

Given a multiset of n positive integers and a target integer t, the Subset Sum problem asks to determine whether there exists a subset of S that sums up to t. The current best deterministic algorithm, by Koiliaris and Xu [SODA'17], runs in $\tilde{O}(\sqrt{n}t)$ time, where $\tilde{O}$ hides polylogarithm factors. Bringmann [SODA'17] later gave a randomized $\tilde{O}(n+t)$ time algorithm using two-stage color-coding. The $\tilde{O}(n+t)$ running time is believed to be near-optimal. In this paper, we present a simple and elegant randomized algorithm for Subset Sum in $\tilde{O}(n+t)$ time. Our new algorithm actually solves its counting version modulo prime $p > t$, by manipulating generating functions using FFT.

Ce Jin, Hongxun Wu
Institute for Interdisciplinary Information Sciences
Tsinghua University
jinc16@mails.tsinghua.edu.cn, wuhx18@mails.tsinghua.edu.cn

CP41

Compressed Sensing with Adversarial Sparse Noise via L1 Regression

We present a simple and effective algorithm for the problem of *sparse robust linear regression*. In this problem, one would like to estimate a sparse vector $w^* \in^n$ from linear measurements corrupted by sparse noise that can arbitrar-

ily change an adversarially chosen η fraction of measured responses y, as well as introduce bounded norm noise to the responses. For Gaussian measurements, we show that a simple algorithm based on L1 regression can successfully estimate w^* for any $\eta < \eta_0 \approx 0.239$, and that this threshold is tight for the algorithm. The number of measurements required by the algorithm is $O(k \log \frac{n}{k})$ for k-sparse estimation, which is within constant factors of the number needed without any sparse noise. Of the three properties we show—the ability to estimate sparse, as well as dense, w^*; the tolerance of a large constant fraction of outliers; and tolerance of adversarial rather than distributional (e.g., Gaussian) dense noise—to the best of our knowledge, no previous polynomial time algorithm was known to achieve more than two.

Sushrut Karmalkar
The University of Texas at Austin
s.sushrut@gmail.com

Eric Price
The Unviersity of Texas at Austin
ecprice@cs.utexas.edu

CP41

Submodular Optimization in the MapReduce Model

Submodular optimization has received significant attention in both practice and theory, as a wide array of problems in machine learning, auction theory, and combinatorial optimization have submodular structure. In practice, these problems often involve large amounts of data, and must be solved in a distributed way. One popular framework for running such distributed algorithms is MapReduce. In this paper, we present two simple algorithms for cardinality constrained submodular optimization in the MapReduce model: the first is a $(1/2 - o(1))$-approximation in 2 MapReduce rounds, and the second is a $(1 - 1/e - \epsilon)$-approximation in $\frac{1+o(1)}{\epsilon}$ MapReduce rounds.

Paul Liu, Jan Vondrak
Stanford University
paul.liu@stanford.edu, jvondrak@stanford.edu

CP41

Simple Contention Resolution via Multiplicative Weight Updates

We consider the classic contention resolution problem, in which devices conspire to share some common resource, for which they each need temporary and exclusive access. To ground the discussion, suppose identical devices wake up at various times, and must send a single packet over a shared multiple-access channel. In each time step they may attempt to send their packet; they receive ternary feedback $\{0, 1, 2^+\}$ from the channel. We prove that a simple strategy suffices to achieve a channel utilization rate of $1/e - O(\epsilon)$, for any $\epsilon > 0$. In each step, device i attempts to send its packet with probability p_i, then applies a rudimentary multiplicative weight-type update to p_i.

$$p_i \leftarrow \begin{cases} p_i \cdot e^{\epsilon} & \text{upon hearing silence } (0) \\ p_i & \text{upon hearing success } (1) \\ p_i \cdot e^{-\epsilon/(e-2)} & \text{upon hearing noise } (2^+) \end{cases}$$

This scheme works well even if the introduction of devices/packets is adversarial, and even if the adversary can *jam* time slots (make noise) at will. If the adversary jams J time slots, then this scheme will achieve channel utilization $1/e - \epsilon$, excluding $O(J)$ wasted slots. Results similar to these (Bender et al. 2016) were already achieved, but with much lower efficiency and a more complex algorithm.

Yi-Jun Chang, Wenyu Jin, Seth Pettie
University of Michigan
cyijun@umich.edu, wyjin@umich.edu, seth@pettie.net

CP41

Approximating Maximin Share Allocations

We study the problem of fair allocation of M indivisible items among N agents using the popular notion of maximin share as our measure of fairness. The maximin share of an agent is the largest value she can guarantee herself if she is allowed to choose a partition of the items into N bundles (one for each agent), on the condition that she receives her least preferred bundle. A maximin share allocation provides each agent a bundle worth at least their maximin share. While it is known that such an allocation need not exist [Procaccia and Wang, 2014, Kurokawa et al., 2016], a series of work [Procaccia and Wang, 2014, Kurokawa et al., 2018, Amanatidis et al., 2017, Barman and Murthy, 2017] provided 2/3 approximation algorithms in which each agent receives a bundle worth at least 2/3 times their maximin share. Recently, [Ghodsi et al., 2018] improved the approximation guarantee to 3/4. Prior works utilize intricate algorithms, with an exception of [Barman and Murthy, 2017] which is a simple greedy solution but relies on sophisticated analysis techniques. In this paper, we propose an alternative 2/3 maximin share approximation which offers both a simple algorithm and straightforward analysis. In contrast to other algorithms, our approach allows for a simple and intuitive understanding of why it works.

Jugal Garg, Peter McGlaughlin
University of Illinois at Urbana-Champaign
jugal@illinois.edu, mcglghl2@illinois.edu

Setareh Taki
University of Illinois at Urbana Champaign
staki2@illinois.edu

CP42

Non-Empty Bins with Simple Tabulation Hashing

We consider the hashing of a set $X \subseteq U$ with $|X| = m$ using a simple tabulation hash function $h : U \to [n] = \{0, \ldots, n-1\}$ and analyze the number of non-empty bins, that is, the size of $h(X)$. We show that the expected size of $h(X)$ matches that with fully random hashing to within low-order terms. We also provide concentration bounds. The number of non-empty bins is a fundamental measure in the balls and bins paradigm, and it is critical in applications such as Bloom filters and Filter hashing. For example, normally Bloom filters are proportioned for a desired low false-positive probability assuming fully random hashing (see en.wikipedia.org/wiki/Bloom_filter). Our results imply that if we implement the hashing with simple tabulation, we obtain the same low false-positive probability for any possible input.

Anders Aamand
University of Copenhagen, BARC
andersaamanda@gmail.com

Mikkel Thorup
University of Copenhagen
mikkel2thorup@gmail.com

CP42

Derandomized Balanced Allocation

In this paper, we study the maximum loads of explicit hash families in the d-choice schemes when allocating sequentially n balls into n bins. We consider the *Uniform-Greedy* scheme, which provides d independent bins for each ball and places the ball into the bin with the least load, and its non-uniform variant — the *Always-Go-Left* scheme introduced by Vöcking. We construct a hash family with $O(\log n \log\log n)$ random bits based on the previous work of Celis et al. and show the following results.

1. With high probability, this hash family has a maximum load of $\frac{\log\log n}{\log d} + O(1)$ in the *Uniform-Greedy* scheme.
2. With high probability, it has a maximum load of $\frac{\log\log n}{d\log\phi_d} + O(1)$ in the *Always-Go-Left* scheme for a constant $\phi_d > 1.61$.

The maximum loads of our hash family match the maximum loads of a perfectly random hash function in the *Uniform-Greedy* and *Always-Go-Left* scheme separately, up to the low order term of constants. Previously, the best known hash families matching the same maximum loads of a perfectly random hash function in d-choice schemes were $O(\log n)$-wise independent functions by Vöcking, which needs $\Theta(\log^2 n)$ random bits.

Xue Chen
Department of Computer Science
University of Texas at Austin
xuechencs@gmail.com

CP42

Optimal Ball Recycling

Balls-and-bins games have been a wildly successful tool for modeling load balancing problems. In this paper, we study a new scenario, which we call the *ball recycling game*, defined as follows: Throw m balls into n bins i.i.d. according to a given probability distribution $\mathbf{p}$. Then, at each time step, pick a non-empty bin and *recycle* its balls: take the balls from the selected bin and re-throw them according to $\mathbf{p}$. This balls-and-bins game closely models memory-access heuristics in databases. The goal is to have a bin-picking method that maximizes the *recycling rate*, defined to be the expected number of balls recycled per step in the stationary distribution. We study two natural strategies for ball recycling: FB, which greedily picks the bin with the maximum number of balls, and RB, which picks a ball at random and recycles its bin. We show that for general $\mathbf{p}$, RB is $O(1)$-optimal, whereas FB can be pessimal. However, when $\mathbf{p} = \mathbf{u}$, the uniform distribution, FB is optimal to within an additive constant.

Michael A. Bender, Jake Christensen
Stony Brook University
bender@cs.stonybrook.edu, jakec43@gmail.com

Alexander Conway, Martin Farach-Colton
Rutgers University
alexander.conway@rutgers.edu, martin@farach-colton.com

Rob Johnson
VMWare Research Group
rob@cs.stonybrook.edu

Meng-Tsung Tsai
National Chiao Tung University
kasuistry@gmail.com

CP42

A Fourier-Analytic Approach for the Discrepancy of Random Set Systems

One of the prominent open problems in combinatorics is the discrepancy of set systems where each element lies in at most t sets. The Beck-Fiala conjecture suggests that the right bound is $O(\sqrt{t})$, but for three decades the only known bound not depending on the size of set system has been $O(t)$. Arguably we currently lack techniques for breaking that barrier. In this paper we introduce discrepancy bounds based on Fourier analysis. We demonstrate our method on random set systems. Suppose one has n elements and m sets containing each element independently with probability p. We prove that in the regime of $n \geq \Theta(m^2 \log(m))$, the discrepancy is at most 1 with high probability. Previously, a result of Ezra and Lovett gave a bound of $O(1)$ under the stricter assumption that $n \gg m^t$.

Rebecca Hoberg, Thomas Rothvoss
University of Washington
rahoberg@gmail.com, rothvoss@uw.edu

CP42

On the Discrepancy of Random Low Degree Set Systems

Motivated by the celebrated Beck-Fiala conjecture, we consider the random setting where there are n elements and m sets and each element lies in t randomly chosen sets. In this setting, Ezra and Lovett showed an $O((t\log t)^{1/2})$ discrepancy bound in the regime when $n \leq m$ and an $O(1)$ bound when $n \gg m^t$. In this paper, we give a tight $O(\sqrt{t})$ bound for the entire range of n and m, under a mild assumption that $t = \Omega(\log\log m)^2$. The result is based on two steps. First, applying the partial coloring method to the case when $n = m\log^{O(1)} m$ and using the properties of the random set system we show that the overall discrepancy incurred is at most $O(\sqrt{t})$. Second, we reduce the general case to that of $n \leq m\log^{O(1)} m$ using LP duality and a careful counting argument.

Nikhil Bansal
CWI and TU Eindhoven
bansal@gmail.com

Raghu Meka
UCLA
raghuvardhan@gmail.com

CP43

Optimal Lower Bounds for Sketching Graph Cuts

We study the space complexity of sketching cuts and Laplacian quadratic forms of graphs. We show that any data structure which approximately stores the sizes of all cuts in an undirected graph on n vertices up to a $1+\epsilon$ error must use $\Omega(n\log n/\epsilon^2)$ bits of space in the worst case, improving the $\Omega(n/\epsilon^2)$ bound of [Andoni et al., ITCS 2016]

and matching the best known upper bound achieved by spectral sparsifiers [Batson et al., STOC 2009]. Our proof is based on a rigidity phenomenon for cut (and spectral) approximation which may be of independent interest: any two $d-$regular graphs which approximate each other's cuts significantly better than a random graph approximates the complete graph must overlap in a constant fraction of their edges.

Charles A. Carlson, Alexandra Kolla
University of Colorado Boulder
chca0914@colorado.edu, alexandra.kolla@colorado.edu

Nikhil Srivastava
University of California Berkeley
nikhil@math.berkeley.edu

Luca Trevisan
UC Berkeley
luca@berkeley.edu

CP43

Short Cycles via Low-diameter Decompositions

We present improved algorithms for *short cycle decomposition* of a graph. Short cycle decompositions were introduced in the recent work of Chu [**?**], and were used to make progress on several questions in graph sparsification. For all constants $\delta \in (0, 1]$, we give an $O(mn^{\delta})$ time algorithm that, given a graph G, partitions its edges into cycles of length $O(\log n)^{\frac{1}{\delta}}$, with $O(n)$ extra edges not in any cycle. This gives the first subquadratic, in fact almost linear time, algorithm achieving polylogarithmic cycle lengths. We also give an $m \cdot \exp(O(\sqrt{\log n}))$ time algorithm that partitions the edges of a graph into cycles of length $\exp(O(\sqrt{\log n} \log \log n))$, with $O(n)$ extra edges not in any cycle. This improves on the short cycle decomposition algorithms given in Chu [**?**] in terms of all parameters, and is significantly simpler. As a result, we obtain faster algorithms and improved guarantees for several problems in graph sparsification – construction of resistance sparsifiers, graphical spectral sketches, degree preserving sparsifiers, and approximating the effective resistances of all edges.

Yang Liu
Stanford University
yangpatil@gmail.com

Sushant Sachdeva, Zejun Yu
University of Toronto
sachdeva@cs.toronto.edu, zejun.yu@main.utoronto.ca

CP43

Spectral Sparsification of Hypergraphs

For an undirected/directed hypergraph $G = (V, E)$, its Laplacian L_G is defined such that its "quadratic form' $x^{\top} L_G(x)$ captures the cut information of G. In particular, $1_S^{\top} L_G(1_S)$ coincides with the cut size of $S \subseteq V$, where 1_S is the characteristic vector of S. A weighted subgraph H of a hypergraph G on a vertex set V is said to be an ϵ-spectral sparsifier of G if $(1 - \epsilon)x^{\top} L_H(x) \leq x^{\top} L_G(x) \leq (1 + \epsilon)x^{\top} L_H(x)$ holds for every x. In this paper, we present a polynomial-time algorithm that, given an undirected/directed hypergraph G on n vertices, constructs an ϵ-spectral sparsifier of G with $O(n^3 \log n/\epsilon^2)$ hyperedges/hyperarcs. The proposed spectral sparsification can be used to improve the time and space complexities of algorithms for solving problems that involve the quadratic form, such as computing the eigenvalues of L_G, computing the effective resistance between a pair of vertices in G, semi-supervised learning based on L_G, and cut problems on G. In addition, our sparsification result implies that any nonnegative submodular function f with $f(\emptyset) = f(V) = 0$ can be concisely represented by a directed hypergraph. Accordingly, we show that, for any distribution, we can properly and agnostically learn submodular functions $f : 2^V \to [0, 1]$ with $f(\emptyset) = f(V) = 0$, with $O(n^4 \log(n/\epsilon)/\epsilon^4)$ samples.

Tasuku Soma
university of Tokyo
tasuku_soma@mist.i.u-tokyo.ac.jp

Yuichi Yoshida
National Institute of Informatics
Preferred Infrastructure
yyoshida@nii.ac.jp

CP43

Expander Decomposition and Pruning: Faster, Stronger, and Simpler

Expander decomposition is a useful method for graph clustering introduced by [Kannan, Vempala, and Vetta, FOCS'00] with wide applications in theory and in practice. In this paper, we study the parametrized version of expander decomposition, where given a graph G of m edges and a parameter ϕ, we find a partition of the vertices into clusters such that each cluster induces a subgraph of conductance at least ϕ (i.e. an expander) and only $\tilde{O}(\phi)$-fraction of edges have endpoints in different clusters. Our algorithm runs in $\tilde{O}(m/\phi)$ time and is the first near-linear time algorithm when ϕ is at least $(\text{polylog}\, m)^{-1}$, which is the case in most practical settings and theoretical applications. This affirmatively answers the open question noted in [Spielman and Teng, STOC'04] and also in [Koutis and Miller, SPAA'08; Orecchia and Vishnoi, SODA'11; Orecchia, Sachdeva, and Vishnoi, STOC'12]. Previous algorithms either take $\Omega(m^{1+o(1)})$ time or only guarantee that each cluster is contained in some unspecified expander. Our decomposition algorithm is developed from first principles based on local flow techniques and is relatively simple. The techniques can be extended to obtain a significant improvement in an *expander pruning* algorithm, which is the key tool for maintaining the expander decomposition on dynamic graphs.

Thatchaphol Saranurak
Toyota Technological Institute at Chicago
saranurak@ttic.edu

Di Wang
Georgia Institute of Technology
wangd@eecs.berkeley.edu

CP43

Cheeger Inequalities for Submodular Transformations

The Cheeger inequality for undirected graphs, which relates the conductance of an undirected graph and the second smallest eigenvalue of its normalized Laplacian, is a cornerstone of spectral graph theory. The Cheeger inequality has been extended to directed graphs and hypergraphs using normalized Laplacians for those, that are no longer

linear but piecewise linear transformations. In this paper, we introduce the notion of a submodular transformation $F : \{0,1\}^n \to R^m$, which applies m submodular functions to the n-dimensional input vector, and then introduce the notions of its Laplacian and normalized Laplacian. With these notions, we unify and generalize the existing Cheeger inequalities by showing a Cheeger inequality for submodular transformations, which relates the conductance of a submodular transformation and the smallest non-trivial eigenvalue of its normalized Laplacian. This result recovers the Cheeger inequalities for undirected graphs, directed graphs, and hypergraphs, and derives novel Cheeger inequalities for mutual information and directed information. Computing the smallest non-trivial eigenvalue of a normalized Laplacian of a submodular transformation is NP-hard under the small set expansion hypothesis. In this paper, we present a polynomial-time $O(\log n)$-approximation algorithm for the symmetric case, which is tight, and a polynomial-time $O(\log^2 n + \log n \cdot \log m)$-approximation algorithm for the general case.

Yuichi Yoshida
National Institute of Informatics
Preferred Infrastructure
yyoshida@nii.ac.jp

CP44

Improved Topological Approximations by Digitization

Čech complexes are useful simplicial complexes for computing and analyzing topological features of data that lies in Euclidean space. Unfortunately, computing these complexes becomes prohibitively expensive for large-sized data sets even for medium-to-low dimensional data. We present an approximation scheme for $(1+\epsilon)$-approximating the topological information of the Čech complexes for n points in R^d, for $\epsilon \in (0,1]$. Our approximation has a total size of $n\left(\frac{1}{\epsilon}\right)^{O(d)}$ for constant dimension d, improving all the currently available $(1+\epsilon)$-approximation schemes of simplicial filtrations in Euclidean space. Perhaps counter-intuitively, we arrive at our result by adding additional $n\left(\frac{1}{\epsilon}\right)^{O(d)}$ sample points to the input. We achieve a bound that is independent of the spread of the point set by pre-identifying the scales at which the Čech complexes changes and sampling accordingly.

Aruni Choudhary
Freie Universität Berlin
aruni.choudhary@fu-berlin.de

Michael Kerber
University of Technology, Graz, Austria
kerber@tugraz.at

Sharath Raghvendra
Virginia Tech
sharathr@vt.edu

CP44

Computing Height Persistence and Homology Generators in R^3 Efficiently

Recently it has been shown that computing the dimension of the first homology group $\mathsf{H}_1(\mathcal{K})$ of a simplicial 2-complex $\mathcal{K}$ embedded linearly in R^4 is as hard as computing the rank of a sparse $0-1$ matrix. This puts a major roadblock to computing persistence and a homology basis (generators) for complexes embedded in R^4 and beyond in less than quadratic or even near-quadratic time. But, what about dimension three? The question for general simplicial complexes $\mathcal{K}$ linearly embedded in R^3 is not completely settled. No algorithm with a complexity better than that of the matrix multiplication is known for this important case. We show that the persistence for *height functions* on such complexes, hence called *height persistence*, can be computed in $O(n \log n)$ time. This allows us to compute a basis (generators) of $\mathsf{H}_i(\mathcal{K})$, $i = 1, 2$, in $O(n \log n + k)$ time where k is the size of the output. This improves significantly the current best bound of $O(n^\omega)$, ω being the matrix multiplication exponent. We achieve these improved bounds by leveraging recent results on zigzag persistence in computational topology, new observations about Reeb graphs, and some efficient geometric data structures.

Tamal K. Dey
Ohio State University
Columbus, Ohio 43210
dey.8@osu.edu

CP44

Constructive Polynomial Partitioning for Algebraic Curves In ^{3} with Applications

In 2015, Guth proved that for any set of k-dimensional varieties in ^{d} and for any positive integer D, there exists a polynomial of degree at most D whose zero-set divides R^d into open connected "cells,' so that only a small fraction of the given varieties intersect each cell. Guth's result generalized an earlier result of Guth and Katz for points. Guth's proof relies on a variant of the Borsuk-Ulam theorem, and for $k > 0$, it is unknown how to obtain an explicit representation of such a partitioning polynomial and how to construct it efficiently. In particular, it is unknown how to effectively construct such a polynomial for curves (or even lines) in R^3. We present an efficient algorithmic construction for this setting. Given a set of n input curves and a positive integer D, we efficiently construct a decomposition of space into $O(D^3 \log^3 D)$ open cells, each of which meets at most $O(n/D^2)$ curves from the input. As an application, we revisit the problem of eliminating depth cycles among non-vertical pairwise disjoint triangles in 3-space, recently studied by Aronov (2017) and de Berg (2017).

Esther Ezra
Georgia Tech
eezra3@math.gatech.edu

Boris Aronov
New York University
boris.aronov@nyu.edu

Joshua Zahl
University of British Columbia
jzahl@math.ubc.ca

CP44

Full Tilt: Universal Constructors for General Shapes with Uniform External Forces

We investigate the problem of assembling general shapes and patterns in a model in which particles move based on uniform external forces until they encounter an obstacle. In this model, corresponding particles may bond when adjacent with one another. Succinctly, this model considers a 2D grid of "open" and "blocked" spaces, along with a set of slidable polyominoes placed at open locations on the board.

The board may be tilted in any of the 4 cardinal directions, causing all slidable polyominoes to move maximally in the specified direction until blocked. By successively applying a sequence of such tilts, tilt sequences provide a method to reconfigure an initial board configuration so as to assemble a collection of previous separate polyominoes into a larger shape. While previous work within this model of assembly has focused on designing a specific board configuration for the assembly of a specific given shape, we propose the problem of designing *universal configurations* that are capable of constructing a large class of shapes and patterns. We also include a study of the complexity of deciding if a particle within a configuration may be relocated to another position, and deciding if a given configuration may be transformed into a second given configuration. In both cases, we show this problem to be PSPACE-complete, even when movable particles are restricted to 1×1 and 2×2 polyominoes that do not stick to one another.

Jose Balanza-Martinez, David Caballero, Angel Cantu, Luis Garcia, Austin Luchsinger, Rene Reyes
University of Texas - Rio Grande Valley
jose.b.balanza@gmail.com, davidcaballeroe@gmail.com, angel.a.cantu@gmail.com, luis.a.garcia01@utrgv.edu, amluchsinger@gmail.com, rfreyes@outlook.com

Robert T. Schweller
University of Texas - Pan American
schwellerr@gmail.com

Tim Wylie
University of Texas - Rio Grande Valley
timothy.wylie@utrgv.edu

CP44
Hardness of Approximation for Morse Matching

Discrete Morse theory has emerged as a powerful tool for a wide range of computational problems in topology. In this context, discrete Morse theory is used to reduce the problem of computing a topological invariant of an input simplicial complex to computing the same topological invariant of a (significantly smaller) collapsed cell or chain complex. Consequently, devising methods for obtaining gradient vector fields on complexes to reduce the size of the problem instance has become an emerging theme over the last decade. While computing the optimal gradient vector field on a simplicial complex is NP-hard, several heuristics have been observed to compute near-optimal gradient vector fields on a wide variety of datasets. Understanding the theoretical limits of these strategies is therefore a fundamental problem in computational topology. In this paper, we establish hardness of approximation results for maximization and minimization variants of the Morse matching problem. In particular, we show that, for a simplicial complex of dimension $d \geq 3$ with n simplices, it is NP-hard to approximate *Min-Morse matching* within a factor of $O(n^{1-\epsilon})$, for any $\epsilon > 0$. Moreover, we establish hardness of approximation results for *Max-Morse matching* for simplicial complexes of dimension $d \geq 2$, using an L-reduction from *Degree 3 Max-Acyclic Subgraph* to *Max-Morse matching*.

Abhishek J. Rathod, Ulrich Bauer
Technische Universität München
rathod@ma.tum.de, mail@ulrich-bauer.org

CP45
High-Dimensional Robust Mean Estimation in Nearly-linear Time

We study the fundamental problem of high-dimensional mean estimation in a robust model where a constant fraction of the samples are adversarially corrupted. Recent work gave the first polynomial time algorithms for this problem with dimension-independent error guarantees for several families of structured distributions. In this work, we give the first nearly-linear time algorithms for high-dimensional robust mean estimation. We focus on distributions with (i) known covariance and sub-gaussian tails, and (ii) unknown bounded covariance. Given N samples on R^d, an ϵ-fraction of which may be arbitrarily corrupted, our algorithms run in time $\tilde{O}(Nd)/\mathrm{poly}(\epsilon)$ and approximate the true mean within the information-theoretically optimal error, up to constant factors. Previous robust algorithms with comparable error guarantees have running times $\tilde{\Omega}(Nd^2)$. Our algorithms rely on a natural family of SDPs parameterized by our current guess ν for the unknown mean $\mu^\star$. We give a win-win analysis as follows: either the primal SDP yields a good estimate for $\mu^\star$, or the dual SDP yields a better guess ν' closer to $\mu^\star$. We exploit the special structure of the corresponding SDPs to show that they are approximately solvable in nearly-linear time. Our approach is quite general, and we believe it can also be applied to obtain nearly-linear time algorithms for other high-dimensional robust learning problems.

Yu Cheng
Duke University
chycharlie@gmail.com

Ilias Diakonikolas
University of Southern California
ilias.diakonikolas@gmail.com

Rong Ge
Duke University
rongge@cs.duke.edu

CP45
Adaptive Sparse Recovery with Limited Adaptivity

The goal of adaptive sparse recovery is to estimate an approximately sparse vector x from a series of linear measurements $A_1x, A_2x, \ldots, A_Rx$, where each matrix A_i may depend on the previous observations. With an unlimited number of rounds R, it is known that $O(k \log\log n)$ measurements suffice for $O(1)$-approximate k-sparse recovery in $\mathbf{R}^n$, and that $\Omega(k + \log\log n)$ measurements are necessary. We initiate the study of what happens with a small, constant number amount of adaptivity. Previous techniques could not give nontrivial bounds using less than 5 rounds of adaptivity, and were inefficient for any constant R. We give nearly matching upper and lower bounds for any constant number of rounds R. Our lower bound shows that $\Omega(k(\log \frac{n}{k})^{1/R})$ measurements are necessary for any $k < 2^{(\log \frac{n}{k})^{1/R}}$; significantly, this is the first lower bound that combines k and n in an adaptive setting. Our upper bound shows that $O(k(\log \frac{n}{k})^{1/R} \cdot \log^* k)$ measurements suffice. The $O(\log^* k)$ gap between the two bounds comes from a similar gap for nonadaptive sparse recovery in the high-SNR regime, and would be reduced to constant factors with improvements to nonadaptive high-SNR sparse recovery.

Akshay Kamath
The University of Texas at Austin
kamath@cs.utexas.edu

Eric Price
The Unviersity of Texas at Austin
ecprice@cs.utexas.edu

CP45

Efficient Algorithms and Lower Bounds for Robust Linear Regression

We study the prototypical problem of high-dimensional linear regression in a robust model where an ϵ-fraction of the samples can be adversarially corrupted. We focus on the fundamental setting where the covariates of the uncorrupted samples are drawn from a Gaussian distribution $\mathcal{N}(0, \Sigma)$ on R^d. We give nearly tight upper bounds and computational lower bounds for this problem as follows:

- For the case that the covariance matrix is known to be the identity, we give a computationally efficient algorithm that draws $\tilde{O}(d/\epsilon^2)$ labeled examples and approximates the unknown regression vector within ℓ_2-norm $O(\epsilon \log(1/\epsilon)\sigma)$. An error of $\Omega(\epsilon\sigma)$ is information-theoretically necessary. Hence, the error guarantee of our algorithm is optimal, up to a logarithmic factor in $1/\epsilon$.
- For the case of unknown covariance Σ, we show that we can efficiently achieve the same error guarantee as in the known covariance case, using an additional $\tilde{O}(d^2/\epsilon^2)$ unlabeled examples. On the other hand, an error of $O(\epsilon\sigma)$ can be information-theoretically attained with $O(d/\epsilon^2)$ samples. We prove a Statistical Query (SQ) lowerbound showing that any polynomial time SQ algorithm (in Huber's contamination model) with estimation complexity $O(d^{1.99})$, must incur an error of $\Omega(\sqrt{\epsilon}\sigma)$.

Weihao Kong
Stanford University
kweihao@gmail.com

Ilias Diakonikolas
USC
diakonik@usc.edu

Alistair Stewart
Stanford
stewart.al@gmail.com

CP45

Dimension-independent Sparse Fourier Transform

The Discrete Fourier Transform (DFT) is a fundamental computational primitive, and the fastest known algorithm for computing the DFT is the FFT (Fast Fourier Transform) algorithm. One remarkable feature of FFT is the fact that its runtime depends only on the size N of the input vector, but *not on d, the dimensionality of the input domain*: FFT runs in time $O(N \log N)$, where $N = n^d$. The state of the art for Sparse FFT, i.e. the problem of computing the DFT of a signal with at most k nonzeros in Fourier domain, is very different: all current techniques for sublinear Sparse FFT incur an exponential dependence on the dimension d. In this paper, we give the first algorithm that computes the DFT of a k-sparse signal in time $(k, \log N)$ *in any dimension d*, avoiding the curse of dimensionality inherent in all previously known techniques. Our main tool is a new class of filters that we refer to as *adaptive aliasing filters*: these filters allow isolating frequencies of a k-Fourier sparse signal using $O(k)$ samples in time domain and $O(k \log N)$ runtime per frequency, in any dimension d. We also investigate natural average case models of the input signal: **(1)** worst-case support in Fourier domain with randomized coefficients and **(2)** random locations in Fourier domain with worst case coefficients. Our techniques lead to an $\widetilde{O}(k^2)$ time algorithm for the former and an $\widetilde{O}(k)$ time algorithm for the latter.

Michael Kapralov
IBM T. J. Watson Research Center,
michael.kapralov@epfl.ch

Ameya Velingker
Google AI
ameya.velingker@epfl.ch

Amir Zandieh
EPFL
amir.zandieh@epfl.ch

CP45

Relative Error Tensor Low Rank Approximation

We consider relative error low rank approximation of tensors: given an order-q tensor $A \in \mathbb{R}^{\Pi_{i=1}^q n_i}$, output a rank-$k$ tensor B for which $\|A - B\|_F^2 \leq (1+\epsilon)\mathrm{OPT}$, where $\mathrm{OPT}= \inf_{\text{rank-}k\ A'} \|A - A'\|_F^2$. One structural issue for obtaining relative error low rank approximation is that there may be no rank-k tensor A_k achieving the above infinum. Another, computational issue, is that an efficient relative error low rank approximation for tensors would allow one to compute the rank of a tensor, which is NP-hard. We bypass these issues via (1) bicriteria and (2) parameterized complexity solutions: (1) We give an algorithm which outputs a rank $k' = O((k/\epsilon)^{q-1})$ tensor B for which $\|A-B\|_F^2 \leq (1+\epsilon)\mathrm{OPT}$ in $nnz(A)+n\cdot\mathrm{poly}(k/\epsilon)$ time in the real RAM model. Here $nnz(A)$ is the number of non-zero entries in A. (2) We give an algorithm for any $\delta > 0$ which outputs a rank k tensor B for which $\|A-B\|_F^2 \leq (1+\epsilon)\mathrm{OPT}$ and runs in $(nnz(A)+n\cdot\mathrm{poly}(k/\epsilon)+\exp(k^2/\epsilon))\cdot n^\delta$ time in the unit cost RAM model. For outputting a rank-k tensor, or even a bicriteria solution, we show a $2^{\Omega(k^{1-o(1)})}$ time lower bound under the Exponential Time Hypothesis. We also obtain new results for matrices, such as $nnz(A)$-time CUR decompositions, which may be of independent interest.

Zhao Song
UT-Austin
zhaos@utexas.edu

David Woodruff
Carnegie Mellon University
dwoodruf@cs.cmu.edu

Peilin Zhong
Columbia University
peilin.zhong@columbia.edu

CP46

Stochastic Matching with Few Queries: New Algorithms and Tools

We consider the stochastic matching problem: A graph $G = (V, E)$ along with a parameter $p \in (0, 1)$ is given in the input. Each edge of G is realized independently with probability p. The goal is to select a degree bounded subgraph H of G such that the expected maximum match-

ing size/weight of H is close to that of G. This model of stochastic matching has attracted significant attention over the recent years. The fundamental open question is the best approximation factor achievable for such algorithms. Prior work has identified breaking (near) half-approximation as a barrier for both weighted and unweighted graphs. We give an algorithm that for unweighted graphs, finds a 0.6568 approximation by querying $\tilde{O}(1/p)$ edges per vertex. This is a significant improvement over the known close-to half approximations of the literature. In particular, it improves the state-of-the-art 0.5001 approximation of Assadi et al. [EC'17]. We also show that the the same algorithm achieves a 0.501 approximation for weighted graphs by querying $\tilde{O}(1/p)$ edges per vertex. This improves both, the approximation factor and the per-vertex queries of the algorithms by Yamaguchi and Maehara [SODA'18] and Behnezhad and Reyhani [EC'18]. Our algorithms are fundamentally different from those considered in the literature, yet are very simple and natural. We consider the new algorithms and our analytical tools to be the main contributions of this paper.

Soheil Behnezhad, Alireza Farhadi
University of Maryland
soheil@cs.umd.edu, farhadi94@gmail.com

MohammadTaghi Hajiaghayi
University of Maryland, College Park
hajiagha@cs.umd.edu

Nima Reyhani
University of Maryland
nima.reyhani@gmail.com

CP46

Popular Matchings and Limits to Tractability

We consider *popular matching* problems in both bipartite and non-bipartite graphs with strict preference lists. It is known that every stable matching is a min-size popular matching. A subclass of max-size popular matchings called *dominant matchings* has been well-studied in bipartite graphs: they always exist and there is a simple linear time algorithm to find one. We show that stable and dominant matchings are the only two tractable subclasses of popular matchings in bipartite graphs; more precisely, we show that it is NP-complete to decide if G admits a popular matching that is neither stable nor dominant. We also show a number of related hardness results, such as (tight) inapproximability of the maximum weight popular matching problem. In non-bipartite graphs, we show a strong negative result: it is NP-hard to decide whether a popular matching exists or not, and the same result holds if we replace *popular* with *dominant*. On the positive side, we show the following results in any graph:

- we identify a subclass of dominant matchings called *strongly dominant* matchings and show a linear time algorithm to decide if a strongly dominant matching exists or not;
- we show an efficient algorithm to compute a popular matching of minimum cost in a graph with edge costs and bounded treewidth.

Yuri Faenza
Ecole Polytechnique Federale de Lausanne
yf2414@columbia.edu

Telikepalli Kavitha
Tata Institute of Fundamental Research
kavitha@tcs.tifr.res.in

Vladlena Powers, Xingyu Zhang
Columbia University
vp2324@columbia.edu, xz2464@columbia.edu

CP46

Beating Greedy for Stochastic Bipartite Matching

We consider the maximum bipartite matching problem in stochastic settings, namely the query-commit and price-of-information models. In the query-commit model, an edge e independently exists with probability p_e. We can query whether an edge exists or not, but if it does exist, then we have to take it into our solution. In the unweighted case, one can query edges in the order given by the classical online algorithm of Karp, Vazirani, and Vazirani to get a $(1-1/e)$-approximation. In contrast, the previously best known algorithm in the weighted case is the $(1/2)$-approximation achieved by the greedy algorithm that sorts the edges according to their weights and queries in that order. Improving upon the basic greedy, we give a $(1-1/e)$-approximation algorithm in the weighted query-commit model. We use a linear program (LP) to upper bound the optimum achieved by any strategy. The proposed LP admits several structural properties that play a crucial role in the design and analysis of our algorithm. We also extend these techniques to get a $(1-1/e)$-approximation algorithm for maximum bipartite matching in the price-of-information model introduced by Singla, who also used the basic greedy algorithm to give a $(1/2)$-approximation.

Buddhima Gamlath, Sagar Kale
École polytechnique fédérale de Lausanne
buddhima.gamlath@epfl.ch, sagar.kale@epfl.ch

Ola Svensson
EPFL
ola.svensson@epfl.ch

CP46

A (1+1/e)-Approximation Algorithm for Maximum Stable Matching with One-sided Ties and Incomplete Lists

We study the problem of finding large weakly stable matchings when preference lists are incomplete and contain one-sided ties. Computing maximum weakly stable matchings is known to be NP-hard. We present a polynomial-time algorithm that achieves an improved approximation ratio of 1+1/e. Building on existing approximation algorithms for this problem, our algorithm is motivated by a proposal process in which numerical priorities are adjusted according to the solution of a linear program, and are used for tie-breaking purposes. Our main idea is to use an infinitesimally small step size for incrementing the priorities. Our analysis involves solving an infinite-dimensional factor-revealing linear program. We also show that the ratio 1+1/e is an upper bound for the integrality gap, which matches the known lower bound.

Chi-Kit Lam, C. Gregory Plaxton
University of Texas at Austin
geocklam@cs.utexas.edu, plaxton@cs.utexas.edu

CP46

Tight Competitive Ratios of Classic Matching Al-

gorithms in the Fully Online Model

Huang et al. (STOC 2018) introduced the fully online matching problem, a generalization of the classic online bipartite matching problem in that it allows all vertices to arrive online and considers general graphs. They showed that the ranking algorithm by Karp et al. (STOC 1990) is strictly better than 0.5-competitive and the problem is strictly harder than the online bipartite matching problem in that no algorithms can be $(1-1/e)$-competitive. This paper pins down two tight competitive ratios of classic algorithms for the fully online matching problem. For the fractional version of the problem, we show that a natural instantiation of the water-filling algorithm is $2-\sqrt{2} \approx 0.585$-competitive, together with a matching hardness result. Interestingly, our hardness result applies to arbitrary algorithms in the edge-arrival models of the online matching problem, improving the state-of-art $\frac{1}{1+\ln 2} \approx 0.5906$ upper bound. For integral algorithms, we show a tight competitive ratio of ≈ 0.567 for the ranking algorithm on bipartite graphs, matching a hardness result by Huang et al. (STOC 2018).

Zhiyi Huang
The University of Hong Kong
zhiyi@cs.hku.hk

Binghui Peng
Tsinghua University
pbh15@mails.tsinghua.edu.cn

Zhihao Gavin Tang
The University of Hong Kong
zhtang@cs.hku.hk

Runzhou Tao
Tsinghua University
trz15@mails.tsinghua.edu.cn

Xiaowei Wu
City University of Hong Kong
wxw0711@gmail.com

Yuhao Zhang
The University of Hong Kong
yhzhang2@cs.hku.hk

CP46

Popular Matching in Roommates Setting is Np-Hard

An input to the Popular Matching problem, in the roommates setting, consists of a graph G and each vertex ranks its neighbors in strict order, known as its preference. In the Popular Matching problem the objective is to test whether there exists a matching $M^\star$ such that there is no matching M where more people are happier with M than with $M^\star$. In this paper we settle the computational complexity of the Popular Matching problem in the roommates setting by showing that the problem is NP-complete. Thus, we resolve an open question that has been repeatedly, explicitly asked over the last decade.

Pranabendu Misra
UiB
pranabendu.m@gmail.com

Sushmita Gupta
University of Bergen
sushmita.gupta@gmail.com

Saket Saurabh
IMSc +UiB
saket@imsc.res.in

Meirav Zehavi
Ben-Gurion University
zehavimeirav@gmail.com

CP47

Approximating (k,l)-Center Clustering for Curves

The Euclidean k-center problem is a classical problem that has been extensively studied in computer science. Given a set $\mathcal{G}$ of n points in Euclidean space, the problem is to determine a set $\mathcal{C}$ of k centers (not necessarily part of $\mathcal{G}$) such that the maximum distance between a point in $\mathcal{G}$ and its nearest neighbor in $\mathcal{C}$ is minimized. In this paper we study the corresponding (k,ℓ)-center problem for polygonal curves under the Fréchet distance, that is, given a set $\mathcal{G}$ of n polygonal curves, each of complexity m, determine a set $\mathcal{C}$ of k polygonal curves, each of complexity ℓ, such that the maximum Fréchet distance of a curve in $\mathcal{G}$ to its closest curve in $\mathcal{C}$ is minimized. In their 2016 paper, Driemel, Krivošija, and Sohler give a near-linear time $(1+\varepsilon)$-approximation algorithm for one-dimensional curves, assuming that k and ℓ are constants. In this paper, we substantially extend and improve the known approximation bounds for curves in dimension 2 and higher. Our approximation bounds are close to being tight.

Kevin Buchin
TU Eindhoven
k.a.buchin@tue.nl

Anne Driemel
University of Bonn
driemel@cs.uni-bonn.de

Joachim Gudmundsson
University of Sydney
joachim.gudmundsson@gmail.com

Michael Horton
New York University
University of Sydney
michael.j.horton@protonmail.com

Irina Kostitsyna
Technical University Eindhoven
i.kostitsyna@tue.nl

Maarten Loffler
Utrecht University
m.loffler@uu.nl

Martijn Struijs
TU Eindhoven
m.a.c.struijs@student.tue.nl

CP47

Exact Algorithms and Lower Bounds for Stable Instances of Euclidean K-Means

We investigate the complexity of solving stable or perturbation-resilient instances of k-means and k-median

clustering in fixed dimension Euclidean metrics (or more generally doubling metrics). The notion of stable or perturbation resilient instances was introduced by Bilu and Linial [2010] and Awasthi, Blum, and Sheffet [2012]. In our context, we say a k-means instance is α-stable if there is a unique optimum solution which remains unchanged if distances are (non-uniformly) stretched by a factor of at most α. Stable clustering instances have been studied to explain why heuristics such as Lloyd's algorithm perform well in practice. In this work we show that for any fixed $\epsilon > 0$, $(1 + \epsilon)$-stable instances of k-means in doubling metrics, which include fixed-dimensional Euclidean metrics, can be solved in polynomial time. We complement this result by showing that under a plausible PCP hypothesis, this is essentially tight: that when the dimension d is part of the input, there is a fixed $\epsilon_0 > 0$ such there is not even a PTAS for $(1+\epsilon_0)$-stable k-means in $\mathbb{R}^d$ unless NP=RP. Given this hypothesis, we consider "stability-preserving" reductions to prove our hardness for stable k-means. Such reductions seem to be more fragile and intricate than standard L-reductions and may be of further use to demonstrate other stable optimization problems are hard to solve.

Zachary Friggstad
University of Alberta, Canada
zacharyf@ualberta.ca

Kamyar Khodamoradi, Mohammad Salavatipour
University of Alberta
khodamor@ualberta.ca, mrs@ualberta.ca

CP47

Frchet Distance under Translation: Conditional Hardness and an Algorithm via Offline Dynamic Grid Reachability

The discrete Fréchet distance is a popular measure for comparing polygonal curves. An important variant is the discrete Fréchet distance under translation, which enables detection of similar movement patterns in different spatial domains. For polygonal curves of length n in the plane, the fastest known algorithm runs in time $\tilde{\mathcal{O}}(n^5)$ [Ben Avraham, Kaplan, Sharir '15]. This is achieved by constructing an arrangement of disks of size $\mathcal{O}(n^4)$, and then traversing its faces while updating reachability in a directed grid graph of size $N := \mathcal{O}(n^2)$, which can be done in time $\tilde{\mathcal{O}}(\sqrt{N})$ per update [Diks, Sankowski '07]. The contribution of this paper is two-fold. First, although it is an open problem to solve dynamic reachability in directed grid graphs faster than $\tilde{\mathcal{O}}(\sqrt{N})$, we improve this part of the algorithm: We observe that an offline variant of dynamic s-t-reachability in directed grid graphs suffices, and we solve this variant in amortized time $\tilde{\mathcal{O}}(N^{1/3})$ per update, resulting in an improved running time of $\tilde{\mathcal{O}}(n^{4.66..})$ for the discrete Fréchet distance under translation. Second, we provide evidence that constructing the arrangement of size $\mathcal{O}(n^4)$ is necessary in the worst case, by proving a conditional lower bound of $n^{4-o(1)}$ on the running time for the discrete Fréchet distance under translation, assuming the Strong Exponential Time Hypothesis.

Karl Bringmann
Max Planck Institute for Informatics,
Saarland Informatics Campus, Germany
kbringma@mpi-inf.mpg.de

Marvin Künnemann
Max-Planck-Institut für Informatik, Saarbrücken,
Germany
marvin@mpi-inf.mpg.de

André Nusser
Max Planck Institute for Informatics
Saarland Informatics Campus
anusser@mpi-inf.mpg.de

CP47

Seth Says: Weak Fréchet Distance is Faster, But Only if it is Continuous and in One Dimension

We show by reduction from the Orthogonal Vectors problem that algorithms with strongly subquadratic running time cannot approximate the Fréchet distance between curves better than a factor 3 unless SETH fails. We show that similar reductions cannot achieve a lower bound with a factor better than 3. Our lower bound holds for the continuous, the discrete, and the weak discrete Fréchet distance even for curves in one dimension. Interestingly, the continuous weak Fréchet distance behaves differently. Our lower bound still holds for curves in two dimensions and higher. However, for curves in one dimension, we provide an exact algorithm to compute the weak Fréchet distance in linear time.

Tim Ophelders
Department of Mathematics and Computer Science
TU Eindhoven
ophelder@egr.msu.edu

Kevin Buchin
TU Eindhoven
k.a.buchin@tue.nl

Bettina Speckmann
Dept. of Mathematics and Computer Science
TU Eindhoven
b.speckmann@tue.nl

CP47

Analysis of Ward's Method

We study *Ward's method* for the hierarchical k-means problem in $\mathbb{R}^d$. This popular greedy heuristic is based on the *complete linkage* paradigm: Starting with all data points as singleton clusters, it successively merges two clusters to form a clustering with one cluster less. The pair of clusters is chosen to (locally) minimize the k-means cost of the clustering in the next step. Complete linkage algorithms are very popular for hierarchical clustering problems, yet their theoretical properties have been studied very little. For the k-center problem, the k-clustering in the hierarchy computed by complete linkage has a worst-case approximation ratio of $\Omega(\log k)$ in general metric spaces, and the best known upper bound is $O(d)$ for inputs in $\mathbb{R}$ with constant d. Complete linkage for k-median/k-means has not been analyzed so far. We show that Ward's method computes a 2-approximation if the optimal k-clustering is well separated. If the optimal clustering additionally satisfies a balance condition, then Ward's method fully recovers the optimum solution. These results hold in arbitrary dimension. We accompany our positive results with a lower bound of $\Omega((3/2)^d)$ for data sets in $\mathbb{R}^d$ that holds if no separation is guaranteed, and with lower bounds when the guaranteed separation is not sufficiently strong. Finally, we show that Ward produces an $O(1)$-approximative clustering when $d = 1$.

Anna Großwendt, Heiko Röglin

University of Bonn
grosswen@cs.uni-bonn.de, heiko@roeglin.org

Melanie Schmidt
TU Dortmund
melanieschmidt@uni-bonn.de

Clemens Rösner
University of Bonn
roesner@cs.uni-bonn.de