Final Program

ACM-SIAM Symposium on Discrete Algorithms

January 5–8, 2020 • Salt Lake City, Utah, U.S.

Symposium on Algorithm Engineering and Experiments

January 5–6, 2020 • Salt Lake City, Utah, U.S.

Symposium on Simplicity in Algorithms

January 6–7, 2020 • Salt Lake City, Utah, U.S.

Symposium on Algorithmic Principles of Computer Systems

January 8, 2020 • Salt Lake City, Utah, U.S.

The Theory Underlying Algorithms Workshop (TUNGA) will take place on January 8, 2020.

APOCS is sponsored by the SIAM Activity Group on Applied and Computational Discrete Algorithms.

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

This activity group 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.

SIAM Events Mobile App

Scan the QR code with any QR reader and download the TripBuilder EventMobile™ app to your iPhone, iPad, iTouch or Android mobile device.
You can also visit http://www.tripbuildermedia.com/apps/siamevents

Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: +1-215-382-9800 Fax: +1-215-386-7999
Conference E-mail: meetings@siam.org • Conference Web: *www.siam.org/meetings/*
Membership and Customer Service:
(800) 447-7426 (U.S. & Canada) or +1-215-382-9800 (worldwide)
https://www.siam.org/conferences/cm/conference/soda20
https://www.siam.org/conferences/cm/conference/alenex20
https://www.siam.org/conferences/cm/conference/sosa20
https://www.siam.org/conferences/cm/conference/apocs20

Table of Contents

SODA COMMITTEES

Program Committee Chair

Shuchi Chawla
University of Wisconsin, Madison, U.S.

Program Committee

Alexandr Andoni
Columbia University, U.S.
Sepehr Assadi
Princeton University and Rutgers University, U.S.
Yossi Azar
Tel Aviv University, Israel
Nikhil Bansal
Centrum Wiskunde and Informatica, Amsterdam, The Netherlands
Hans Bodlaender
Utrecht University, The Netherlands
Mark Bun
Boston University, U.S.
Sergio Cabello
University of Ljubljana, Slovenia
Karthekeyan Chandrasekaran
University of Illinois Urbana-Champaign, U.S.
Shiri Chechik
Tel Aviv University, Israel
Maria Chudnovsky
Princeton University, U.S.
Radu Curticapean
IT University of Copenhagen, Denmark
Michael Dinitz
Johns Hopkins University, U.S.
Anne Driemel
University of Bonn, Germany
Charilaos Efthymiou
Warwick University, United Kingdom
Moran Feldman
Open University of Israel, Israel
Hu Fu
University of British Columbia, Canada
Jugal Garg
University of Illinois Urbana-Champaign, U.S.
Petr Hlineny
Masaryk University, Czech Republic
Sungjin Im
University of California Merced, U.S.
Gautam Kamath
University of Waterloo, Canada
Mihyun Kang
Graz University of Technology, Austria
Michael Kapralov
École Polytechnique Fédérale de Lausanne, Switzerland
Robert Kleinberg
Cornell University, U.S.
Ravishankar Krishnaswamy
Microsoft Research Bangalore, India
Fabian Kuhn
University of Freiburg, Germany
Janardhan Kulkarni
Microsoft Research Redmond, U.S.
Amit Kumar
Indian Institute of Technology Delhi, India
Gabor Kun
Hungarian Academy of Sciences, Hungary
Laszlo Miklos Lovasz
Massachusetts Institute of Technology, U.S.
Konstantin Makarychev
Northwestern University, U.S.
Marco Molinaro
Pontifical Catholic University of Rio de Janeiro, Brazil
Wolfgang Mulzer
Freie Universität Berlin, Germany
Kamesh Munagala
Duke University, U.S.
Jesper Nederlof
Eindhoven University, The Netherlands
Jelani Nelson
Harvard University, U.S.
Huy Le Nguyen
Northeastern University, U.S.
Lorenzo Orecchia
Boston University, U.S.
Yuval Rabani
Hebrew University, Israel
Barna Saha
University of Massachusetts, Amherst, U.S.
Laura Sanita
University of Waterloo, Canada
Melanie Schmidt
University of Bonn, Germany
Bruce Shepherd
University of British Columbia, Canada
Anastasios Sidiropoulos
University of Illinois at Chicago, U.S.
Sahil Singla
Princeton University and Institute for Advanced Study, U.S.
Balasubramanian Sivan
Google, U.S.
Shay Solomon
Tel Aviv University, Israel
Thomas Steinke
IBM Research Almaden, U.S.
Avishay Tal
Stanford University, U.S.
Christos Tzamos
University of Wisconsin-Madison, U.S.
Chris Umans
California Institute of Technology, U.S.
Virginia Vassilevska Williams
Massachusetts Institute of Technology, U.S.
Nisheeth Vishnoi
Yale University, U.S.
David P. Williamson
Cornell University, U.S.
David Woodruff
Carnegie Mellon University, U.S.
Mary Wootters
Stanford University, U.S.

Steering Committee Chair

Shang-Hua Teng
University of Southern California, U.S.

Steering Committee

Julia Chuzhoy
Toyota Technological Institute at Chicago, U.S.
Piotr Indyk
Massachusetts Institute of Technology, U.S.
Pavol Hell
Simon Fraser University, Canada
Daniel Král
University of Warwick, United Kingdom
Cliff Stein
Columbia University, U.S. (ex-officio member)

ALENEX COMMITTEES

Program Committee Co-Chairs

Guy Blelloch
Carnegie Mellon University, U.S.
Irene Finocchi
Sapienza University of Rome, Italy

Program Committee

Rudolf Fleischer
German University of Technology in Oman, Oman
Loukas Georgiadis
University of Ioannina, Greece
Jakub Łacki
Google, U.S.
Hyeontaek Lim
Google, U.S.

Kamesh Madduri
Pennsylvania State University, U.S.
Andrea Marino
University of Florence, Italy
Petra Mutzel
Technical University of Dortmund, Germany
Gonzalo Navarro
University of Chile, Chile
Cynthia Phillips
Sandia National Laboratories, U.S.
Christian Schulz
University of Vienna, Austria
Nodari Sitchinava
University of Hawaii at Manoa, U.S.
Yihan Sun
Carnegie Mellon University, U.S.
Alexander van der Grinten
Humboldt University of Berlin, Germany

Steering Committee Chair

Andrew V. Goldberg
Amazon.com, U.S.

Steering Committee

Giuseppe F. Italiano
University of Rome "Tor Vergata", Italy
Henning Meyerhenke
Humboldt University, Germany
Rasmus Pagh
IT University of Copenhagen, Denmark
Vijaya Ramachandran
University of Texas, Austin, U.S
Shang-Hua Teng
University of Southern California, U.S.
Dorothea Wagner
Karlsruhe Institute of Technology, Germany

SOSA COMMITTEES

Program Committee Co-Chairs

Martin Farach-Colton
Rutgers University, U.S.
Inge Li Gørtz
Technical University of Denmark, Denmark

Program Committee

Kunal Agrawal
Washington University in St. Louis, U.S.
Guy Blelloch
Carnegie Mellon University, U.S.
Vladimir Braverman
Johns Hopkins University, U.S.
Karl Bringmann
Max Planck Institute for Informatics, Germany
Deeparnab Chakrabarty
Dartmouth College, U.S.
Moses Charikar
Stanford University, U.S.
Artur Czumaj
University of Warwick, United Kingdom
David Eppstein
University of California, Irvine, U.S.
Naveen Garg
Indian Institute of Technology Delhi, India
Paweł Gawrychowski
University of Wrocław, Poland
Anupam Gupta
Carnegie Mellon University, U.S.
Rob Johnson
VMware, U.S.
Michal Koucky
Charles University, Czech Republic
Miguel A. Mosteiro
Pace University, U.S.
Danupon Nanongkai
KTH Royal Institute of Technology, Sweden
Vijaya Ramachandran
The University of Texas at Austin, U.S.
R. Ravi
Carnegie Mellon University, U.S.
Andrea Richa
Arizona State University, U.S.
Shikha Singh
Williams College, U.S.
Bettina Speckmann
TU Eindhoven, The Netherlands
Tatiana Starikovskaya
École Normale Supérieure, France
Kunal Talwar
Google, U.S.
Przemek Uznański
University of Wrocław, Poland
Oren Weimann
University of Haifa, Israel

Steering Committee

Michael A. Bender
Stony Brook University, U.S.
David Karger
Massachusetts Institute of Technology, U.S.
Tsvi Kopelowitz
Bar-Ilan University, Israel
Seth Pettie
University of Michigan, U.S.
Robert Tarjan
Princeton University, U.S.
Mikkel Thorup
University of Copenhagen, Denmark

APOCS COMMITTEES

Program Committee Chair

Bruce Maggs
Duke University and Akamai Technologies, U.S.

Program Committee

Rachit Agarwal
Cornell University, U.S.
Michael Dinitz
Johns Hopkins University, U.S.
Phillip Gibbons
Carnegie Mellon University, U.S.
Yan Gu
Massachusetts Institute of Technology, U.S.
Angelina Lee
Washington University in St. Louis, U.S.
Michael Mitzenmacher
Harvard University, U.S.
Cynthia Phillips
Sandia National Laboratories, U.S.
Michael Schapira
Hebrew University of Jerusalem, Israel
Tao Schardl
Massachusetts Institute of Technology, U.S.
Christian Scheideler
University of Paderborn, Germany
Gala Yadgar
Technion Israel Institute of Technology, Israel

Steering Committee Chair

Martin Farach-Colton
Rutgers University U.S.

Steering Committee

Michael Bender
Stony Brook University, U.S.
Guy Blelloch
Carnegie Mellon University, U.S.
Jennifer Chayes
Microsoft Research, U.S.
Charles Leiserson
Massachusetts Institute of Technology, U.S.
Don Porter
University of North Carolina, Chapel Hill, U.S.
Jennifer Rexford
Princeton University, U.S.
Margo Seltzer
University of British Columbia, Canada

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

Located in the Topaz Room, the registration desk is open at the following times:

Saturday, January 4
5:00 PM - 8:00 PM

Sunday, January 5
8:00 AM - 5:00 PM

Monday, January 6
8:00 AM - 5:00 PM

Tuesday, January 7
8:00 AM - 5:00 PM

Wednesday, January 8
8:00 AM - 5:00 PM

Hotel Address

Hilton Salt Lake City Center
255 South West Temple
Salt Lake City, Utah 84101
United States

Hotel Telephone Number

To reach an attendee or leave a message, call +1-801-328-2000. 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

Visit *care.com* for information on child care services. *Care.com* provides a web-based resource to connect individuals with vetted babysitters and nannies. Attendees are responsible for making their own child care arrangements.

Corporate Members and Affiliates

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 for Dynamics of Complex Technical Systems Magdeburg
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 as of November 2019.

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 paid the Non-Member rate to attend, you can apply the difference of $140 between what you paid and what a member paid towards a SIAM membership. Contact SIAM Customer Service for details or join at the conference registration desk.

Students who paid the Student Non-Member Rate will be automatically enrolled as SIAM Student Members. Please go to *https://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 USA

Audio-Visual Set-Up in Meeting Rooms

SIAM does not provide computers. 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 also provide a limited number of email stations.

Registration Fee Includes

Admission to all technical sessions (*SODA, ALENEX, SOSA, APOCS, and TUNGA*)

- Business Meetings
- Coffee breaks daily
- Continental breakfast daily
- Luncheon on Sunday, January 5
- Access to online Proceedings
- Room set-ups and audio-visual equipment
- Welcome Reception

Job Postings

Please check at the registration desk for job postings, or visit *jobs.siam.org/*.

Conference Sponsors

Google

IBM Research

Microsoft

SIGACT

2020 Conference Bag Sponsor

Julia computing

Name Badges

Please complete the emergency contact information on the back of your name badge.

Comments?

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

Social Events

Welcome Reception

Saturday, January 4
6:00 p.m.- 8:00 p.m.
Grand Ballroom A&B

ALENEX Business Meeting

Sunday, January 5
6:45 p.m. - 7:45 p.m.
Canyon A&B

SODA Business Meeting & Awards Presentation, and SOSA Business Meeting

Monday, January 6
6:45 p.m. - 7:45 p.m.
Grand Ballroom C
Complimentary beer and wine will be served.

APOCS Business Meeting

Wednesday, January 8
5:20 p.m.
Canyon A&B

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.

If you have experienced or observed behavior that is not consistent with the principles expressed above, you are encouraged to report any violation using the SIAM hotline, hosted by the third-party hotline provider, EthicsPoint. The information you provide will be sent to us by EthicsPoint on a totally confidential and anonymous basis if you should choose. You have our guarantee that your comments will be heard. Please submit reports at *http://siam.ethicspoint.com/*.

Please Note

SIAM is not responsible for the safety and security of attendees' belongings. Do not leave your property unattended.

Additionally, please silence your devices.

Recording of Presentations

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

Social Media

#SIAMDA20 #ALENEX20
#SOSA20 #APOCS20
@TheSIAMNews

Changes to the Program

Please review the online program schedule (*https://meetings.siam.org/program.cfm?CONFCODE=da20*) or use the mobile app for up-to-date scheduling information.

SIAM Events Mobile App Powered by TripBuilder®

To enhance your conference experience, we're providing a state-of-the-art mobile app to give you important conference information right at your fingertips. With this TripBuilder EventMobile™ app, you can:

- Create your own custom schedule
- View sessions, speakers, exhibitors and more
- Take notes and export them to your email
- View award-winning TripBuilder recommendations for the meeting location
- Get instant alerts about important conference info

Scan the QR code with any QR reader and download the TripBuilder EventMobile™ app to your iPhone, iPad, iTouch or Android mobile device. You can also visit *http://www.tripbuildermedia.com/apps/siamevents*

SIAM Activity Group on Discrete Mathematics (SIAG/DM)

www.siam.org/Activity-Groups/discrete-mathematics

A great way to get involved!

Collaborate and interact with mathematicians and applied scientists whose work involves discrete mathematics.

Society for Industrial and Applied Mathematics

ACTIVITIES INCLUDE

- Special sessions at SIAM Annual Meetings
- Biennial conference on Discrete Mathematics
- Co-sponsors the annual ACM-SIAM Symposium on Discrete Algorithms
- Dénes König Prize
- DM-Net

BENEFITS OF SIAG/DM MEMBERSHIP

- Additional $15 discount on registration at SIAM Conference on Discrete Mathematics (excludes student)
- Electronic communications from your peers about recent developments in your specialty
- Eligibility for candidacy for SIAG/DM office
- Participation in the selection of SIAG/DM officers

ELIGIBILITY

- Be a current SIAM member.

COST

- $15 per year
- Student members can join two activity groups for free!

2018–2020 SIAG/DM OFFICERS

Chair: Bojan Mohar, *Simon Frasier University*
Vice Chair: Blair Sullivan, *North Carolina State University*
Program Director: Seshadhri Comandur, *University of California, Santa Cruz*
Secretary: Richard Brewster, *Thompson Rivers University*

SIAG/DM: ***www.siam.org/Activity-Groups/discrete-mathematics***
SIAM: *siam.org/joinsiam*

Best Paper and Best Student Paper Awards

Best Paper Award

The Best Paper Award is awarded jointly to the following two papers:

Chasing Convex Bodies Optimally

Mark Sellke, *Stanford University, U.S.*

Chasing Convex Bodies with Linear Competitive Ratio

C.J. Argue, Anupam Gupta, *Carnegie Mellon University, U.S.*
Guru Guruganesh, *Google Research, U.S.*
Ziye Tang, *Carnegie Mellon University, U.S.*

These papers will be presented on Tuesday, January 7, in CP26 SODA Session 7A.

See page 22 for session details.

Best Student Paper Award

Chasing Convex Bodies Optimally

Mark Sellke, *Stanford University, U.S.*

This paper will be presented on Tuesday, January 7, in CP26 SODA Session 7A.

See page 22 for session details.

Invited Plenary Speakers

All Invited Plenary Presentations will take place in Grand Ballroom C.

Sunday, January 5

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

IP1 Algorithmic Questions in Machine Learning
Sanjeev Arora, *Princeton University, U.S.*

Monday, January 6

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

IP2 Dynamic Algorithms for Point Sets and Graphs
Monika Henzinger, *University of Vienna, Austria*

Tuesday, January 7

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

IP3 Computing Maximum Independent Set in Graph Classes
Stéphan Thomassé, *École Normale Supérieure de Lyon, France*

6:45 p.m. - 7:45 p.m.

IP4 SIAG/ACDA Plenary Presentation - Models and Algorithms in Support of Modern Urban Mobility
David B. Shmoys, *Cornell University, U.S.*

Wednesday, January 8

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

IP5 Phase Transitions in Random Constraint Satisfaction Problems
Nike Sun, *Massachusetts Institute of Technology, U.S.*

Program Schedule

siam® 2020

ACM-SIAM Symposium on
Discrete Algorithms
January 5–8, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Algorithm Engineering and Experiments
January 5–6, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Simplicity in Algorithms
January 6–7, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Algorithmic Principles of Computer Systems
January 8, 2020 • Salt Lake City, Utah, U.S.

Saturday, January 4

Registration

5:00 p.m.-8:00 p.m.

Room: Topaz

Welcome Reception

6:00 p.m.-8:00 p.m.

Room: Grand Ballroom A&B

Sunday, January 5

Registration

8:00 a.m.-5:00 p.m.

Room: Topaz

Continental Breakfast

8:30 a.m.-9:00 a.m.

Room: Grand Ballroom Foyer

Sunday, January 5

CP1

ALENEX Session 1

9:00 a.m.-10:40 a.m.

Room: Canyon A&B

Chair: Jakub Lacki, Google Research, U.S.

9:00-9:20 Computing Optimal Hypertree Decompositions

Andre Schidler, Technische Universitaet Wien, Austria; Stefan Szeider, Technische Universität Wien, Austria

9:25-9:45 FixCon: A Generic Solver for Fixed-Cardinality Subgraph Problems

Christian Komusiewicz and *Frank Sommer*, Philipps-Universität, Marburg, Germany

9:50-10:10 Engineering Kernelization for Maximum Cut

Damir Ferizovic, Demian Hespe, and *Sebastian Lamm*, Karlsruhe Institute of Technology, Germany; Matthias Mnich, Maastricht University, Netherlands and Universität Bonn, Germany; Christian Schulz, University of Vienna, Austria and Karlsruhe Institute of Technology, Germany; Darren Strash, Colgate University, U.S.

10:15-10:35 Shared-Memory Branch-and-Reduce for Multiterminal Cuts

Monika Henzinger and *Alexander Noe*, University of Vienna, Austria; Christian Schulz, University of Vienna, Austria and Karlsruhe Institute of Technology, Germany

Sunday, January 5

CP2

SODA Session 1A

9:00 a.m.-11:05 a.m.

Room: Grand Ballroom C

Chair: Michael Dinitz, Johns Hopkins University, U.S.

9:00-9:20 An Almost 2-Approximation for All-Pairs of Shortest Paths in Subquadratic Time

Liam Roditty and *Maor Akav*, Bar-Ilan University, Israel

9:25-9:45 Truly Subcubic Min-Plus Product for Less Structured Matrices, with Applications

Yinzhan Xu and Virginia Williams, Massachusetts Institute of Technology, U.S.

9:50-10:10 Equivalences Between Triangle and Range Query Problems

Lech Duraj, Krzysztof Kleiner, and *Adam Polak*, Jagiellonian University, Poland; Virginia Vassilevska Williams, Massachusetts Institute of Technology, U.S.

10:15-10:35 New Algorithms and Lower Bounds for All-Pairs Max-Flow in Undirected Graphs

Ohad Trabelsi, Weizmann Institute of Science, Israel; Amir Abboud, IBM Almaden Research Center, U.S.; Robert Krauthgamer, Weizmann Institute of Science, Israel

10:40-11:00 Tight Running Time Lower Bounds for Strong Inapproximability of Maximum K-Coverage, Unique Set Cover and Related Problems (via T-Wise Agreement Testing Theorem)

Pasin Manurangsi, University of California, Berkeley, U.S.

Sunday, January 5

CP3

SODA Session 1B

9:00 a.m.-11:05 a.m.

Room: Alpine East

Chair: Piotr Indyk, Massachusetts Institute of Technology, U.S.

9:00-9:20 Learning from Satisfying Assignments under Continuous Distributions

Clément Canonne, IBM Research, U.S.; Anindya De, University of Pennsylvania, U.S.; Rocco A. Servedio, Columbia University, U.S.

9:25-9:45 A Tight Analysis of Greedy Yields Subexponential Time Approximation for Uniform Decision Tree

Ray Li, Percy Liang, and Steve Mussmann, Stanford University, U.S.

9:50-10:10 Finding a Latent k–simplex in O*(k • Nnz (data)) Time Via Subset Smoothing

Chiranjib Bhattacharyya, Indian Institute of Science, Bangalore, India; *Ravindran Kannan*, Microsoft Research, India

10:15-10:35 Oblivious Sketching of High-Degree Polynomial Kernels

Thomas Dybdahl Ahle, IT University of Copenhagen, Denmark; Michael Kapralov, École Polytechnique Fédérale de Lausanne, Switzerland; Jakob B. T. Knudsen, University of Copenhagen, Denmark; Rasmus Pagh, IT University of Copenhagen, Denmark; Ameya Velingker, Google, Inc., U.S.; David Woodruff, Carnegie Mellon University, U.S.; *Amir Zandieh*, École Polytechnique Fédérale de Lausanne, Switzerland

10:40-11:00 List Decodable Learning Via Sum of Squares

Morris M. Yau and Prasad Raghavendra, University of California, Berkeley, U.S.

Sunday, January 5

CP4

SODA Session 1C

9:00 a.m.-11:05 a.m.

Room: Alpine West

Chair: To Be Determined

9:00-9:20 Zeros of Ferromagnetic Two-Spin Systems

Heng Guo, University of Edinburgh, United Kingdom; *Jingcheng Liu*, California Institute of Technology, U.S.; Pinyan Lu, Shanghai University of Finance and Economics, China

9:25-9:45 Adaptive Quantum Simulated Annealing for Bayesian Inference and Estimating Partition Functions

Aram Harrow and *Annie Y. Wei*, Massachusetts Institute of Technology, U.S.

9:50-10:10 Optimal Space-Depth Trade-Off of Cnot Circuits in Quantum Logic Synthesis

Jiaqing Jiang and Xiaoming Sun, Chinese Academy of Sciences, China; Shang-Hua Teng, University of Southern California, U.S.; Bujiao Wu, Chinese Academy of Sciences, China; Kewen Wu, Peking University, China; Jialin Zhang, Chinese Academy of Sciences, China

10:15-10:35 Normalizers and Permutational Isomorphisms in Simply-Exponential Time

Daniel Wiebking, RWTH-Aachen, Germany

10:40-11:00 The Directed Flat Wall Theorem

Archontia Giannopoulou, National and Kapodistrian University of Athens, Greece; *Ken-ichi Kawarabayashi*, National Institute of Informatics, Japan; Stephan Kreutzer, Technische Universität Berlin, Germany; O-Joung Kwon, Incheon National University, Korea

Coffee Break

11:05 a.m.-11:30 a.m.

Room: Grand Ballroom Foyer

Sunday, January 5

IP1

Algorithmic Questions in Machine Learning

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

Room: Grand Ballroom C

Chair: Shuchi Chawla, University of Wisconsin, Madison, U.S.

Machine learning (ML) is a very algorithmic discipline. But to experts in traditional algorithms (e.g. from STOC/FOCS/SODA) it is not always clear how to work on ML problems ---especially deep learning--- due to a sizeable gulf in mindsets as well as well-known challenges in thinking about nonconvex optimization. This talk will use a series of vignettes to illustrate interesting theory questions in contemporary ML, and how algorithmic insights have been obtained in recent years.

Sanjeev Arora
Princeton University, U.S.

Luncheon **Ticketed Event**

12:30 p.m.-2:00 p.m.

Room: Grand Ballroom A&B

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

Sunday, January 5

CP5

ALENEX Session 2

2:00 p.m.-3:40 p.m.

Room: Canyon A&B

Chair: Petra Mutzel, Technische Universität Dortmund, Germany

2:00-2:20 Group Centrality Maximization for Large-Scale Graphs

Eugenio Angriman and *Alexander Van Der Grinten*, Humboldt University at Berlin, Germany; Aleksandar Bojchevski, Daniel Zügner, and Stephan Günnemann, Technical University of Munich, Germany; Henning Meyerhenke, Humboldt University at Berlin, Germany

2:25-2:45 Approximating Vertex Cover Using Structural Rounding

Andrew van Der Poel, North Carolina State University, U.S.; *Brian Lavallee*, University of Utah, U.S.; Hayley Russell, North Carolina State University, U.S.; Blair Sullivan, University of Utah, U.S.

2:50-3:10 A Multi-Criteria Approximation Algorithm for Influence Maximization with Probabilistic Guarantees

Nguyen D. Pham and Gopal Pandurangan, University of Houston, U.S.; Anil Vullikanti, University of Virginia, U.S.; Maleq Khan, Texas A&M University, Kingsville, U.S.

3:15-3:35 Puzzling Grid Embeddings

Moritz Beck and Sabine Storandt, Universität Konstanz, Germany

Sunday, January 5

CP6

SODA Session 2A

2:00 p.m.-4:05 p.m.

Room: Grand Ballroom C

Chair: Janardhan Kulkarni, Microsoft Research, U.S.

2:00-2:20 A Deterministic Linear Program Solver in Current Matrix Multiplication Time

Jan Van Den Brand, KTH Royal Institute of Technology, Sweden

2:25-2:45 Packing Lps Are Hard to Solve Accurately, Assuming Linear Equations Are Hard

Rasmus Kyng, ETH Zürich, Switzerland; Di Wang, Google Research, U.S.; *Peng Zhang*, Yale University, U.S.

2:50-3:10 Symmetric Polymorphisms and Efficient Decidability of Promise Csps

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

3:15-3:35 Extended Formulation Lower Bounds for Refuting Random Csps

Jonah Brown-Cohen, KTH Royal Institute of Technology, Sweden; Prasad Raghavendra, University of California, Berkeley, U.S.

3:40-4:00 Quasi-Popular Matchings, Optimality, and Extended Formulations

Yuri Faenza, Columbia University, U.S.; Kavitha Telikepalli, Tata Institute of Fundamental Research, India

Sunday, January 5

CP7

SODA Session 2B

2:00 p.m.-4:05 p.m.

Room: Alpine East

Chair: Piotr Indyk, Massachusetts Institute of Technology, U.S.

2:00-2:20 Computational Concentration of Measure: Optimal Bounds, Reductions, and More

Omid Etesami, Institute for Research in Fundamental Sciences, Iran; *Saeed Mahloujifar* and Mohammad Mahmoody, University of Virginia, U.S.

2:25-2:45 Efficiency of the Floating Body As a Robust Measure of Dispersion

Joseph T. Anderson, Salisbury University, U.S.; Luis Rademacher, University of California, Davis, U.S.

2:50-3:10 Sample Efficient Toeplitz Covariance Estimation

Yonina Eldar, Weizmann Institute of Science, Israel; Jerry Li, Microsoft Research, U.S.; Cameron Musco, University of Massachusetts, Amherst, U.S.; *Christopher Musco*, New York University, U.S.

3:15-3:35 On the Learnability of Random Deep Networks

Rina Panigrahy, Google Research, U.S.

3:40-4:00 Exact Computation of a Manifold Metric, Via Lipschitz Extensions and Shortest Paths on a Graph

Gary Miller, Carnegie Mellon University, U.S.; *Timothy Chu*, Carnegie Mellon University, U.S.; Donald Sheehy, University of Connecticut, U.S.

Sunday, January 5

CP8

SODA Session 2C

2:00 p.m.-4:05 p.m.

Room: Alpine West

Chair: Sergio Cabello, University of Ljubljana, Slovenia

2:00-2:20 How to Store a Random Walk

Emanuele Viola, Northeastern University, U.S.; Omri Weinstein, Columbia University, U.S.; *Huacheng Yu*, Princeton University, U.S.

2:25-2:45 Shorter Labeling Schemes for Planar Graphs

Cyril Gavoille, University of Bordeaux, France; Marthe Bonamy, LABRI, Univ Bordeaux, France; Michal Pilipczuk, University of Warsaw, Poland

2:50-3:10 Dynamic Low-Stretch Spanning Trees in Subpolynomial Time

Shiri Chechik, Tel Aviv University, Israel; *Tianyi Zhang*, Tsinghua University, China

3:15-3:35 Coarse-Grained Complexity for Dynamic Algorithms

Sayan Bhattacharya, University of Warwick, United Kingdom; Danupon Nanongkai, KTH Royal Institute of Technology, Sweden; *Thatchaphol Saranurak*, Toyota Technological Institute at Chicago, U.S.

3:40-4:00 Extremal Distances in Directed Graphs: Tight Spanners and Near-Optimal Approximation Algorithms

Keerti Choudhary, Weizmann Institute of Science, Israel; Omer Gold, Tel Aviv University, Israel

Coffee Break

4:05 p.m.-4:30 p.m.

Room: Grand Ballroom Foyer

Sunday, January 5

CP9

ALENEX Session 3

4:30 p.m.-6:10 p.m.

Room: Canyon A&B

Chair: Alexander Van Der Grinten, Humboldt University at Berlin, Germany

4:30-4:50 Fully Dynamic Single-Source Reachability in Practice: An Experimental Study

Kathrin Hanauer and Monika Henzinger, University of Vienna, Austria; Christian Schulz, University of Vienna, Austria and Karlsruhe Institute of Technology, Germany

4:55-5:15 Approximating Multiobjective Shortest Path in Practice

Fritz Bökler and Markus Chimani, University of Osnabrueck, Germany

5:20-5:40 Karp-Sipser Based Kernels for Bipartite Graph Matching

Ioannis Panagiotas, Inria and ENS Lyon, France; Kamer Kaya, Sabanci University, Turkey; Johannes Langguth, Simula Research Laboratory, Norway; Bora Uçar, CNRS and ENS Lyon, France

5:45-6:05 Shrinking Trees Not Blossoms:A Recursive Maximum Matching Approach

Andre Droschinsky, Petra Mutzel, and *Erik Thordsen*, Technical University of Dortmund, Germany

Sunday, January 5

CP10

SODA Session 3A

4:30 p.m.-6:35 p.m.

Room: Grand Ballroom C

Chair: Janardhan Kulkarni, Microsoft Research, U.S.

4:30-4:50 Exponential Separations in Local Differential Privacy

Matthew Joseph, University of Pennsylvania, U.S.; Jieming Mao, Google Research, U.S.; Aaron Roth, University of Pennsylvania, U.S.

4:55-5:15 Individual Sensitivity Preprocessing for Data Privacy

Rachel Cummings and David Durfee, Georgia Institute of Technology, U.S.

5:20-5:40 Locally Private K-Means Clustering

Uri Stemmer, Ben-Gurion University, Israel

5:45-6:05 Differentially Private Release of Synthetic Graphs

Marek Elias and Michael Kapralov, École Polytechnique Fédérale de Lausanne, Switzerland; Janardhan Kulkarni, Microsoft Research, U.S.; Yin Tat Lee, University of Washington and Microsoft Research, U.S.

6:10-6:30 The Rank of Sparse Random Matrices

Amin Coja-Oghlan, Goethe University, Germany; Alperen Ergür, Carnegie Mellon University, U.S.; Pu Gao, University of Waterloo, Canada; Samuel Hetterich and *Maurice Rolvien*, Goethe University, Germany

Sunday, January 5

CP11

SODA Session 3B

4:30 p.m.-6:35 p.m.

Room: Alpine East

Chair: Sergio Cabello, University of Ljubljana, Slovenia

4:30-4:50 A Lower Bound for Jumbled Indexing

Peyman Afshani, Aarhus University, Denmark; Ingo Van Duijn, Aalborg University, Denmark; *Rasmus Killmann*, Aarhus University, Denmark; Jesper Sindahl Nielsen, Uber, U.S.

4:55-5:15 Locally Consistent Parsing for Text Indexing in Small Space

Shay Golan, Or Birenzwige, and Ely Porat, Bar-Ilan University, Israel

5:20-5:40 Better Data Structures for Colored Orthogonal Range Reporting

Timothy M. Chan, University of Illinois at Urbana-Champaign, U.S.; *Yakov Nekrich*, Michigan Technological University, U.S.

5:45-6:05 Faster Deterministic and Las Vegas Algorithms for Offline Approximate Nearest Neighbors in High Dimensions

Josh Alman, Harvard University, U.S.; Timothy M. Chan, University of Illinois at Urbana-Champaign, U.S.; Ryan Williams, Massachusetts Institute of Technology, U.S.

6:10-6:30 Flushing Without Cascades

Michael A. Bender and Rathish Das, Stony Brook University, U.S.; Martin Farach-Colton, Rutgers University, U.S.; Rob Johnson, Vmware Research, U.S.; *William Kuszmaul*, Massachusetts Institute of Technology, U.S.

Sunday, January 5

CP12

SODA Session 3C

4:30 p.m.-6:35 p.m.

Room: Alpine West

Chair: Artur Czumaj, University of Warwick, United Kingdom

4:30-4:50 A Randomly Weighted Minimum Spanning Tree with a Random Cost Constraint

Alan Frieze, Carnegie Mellon University, U.S.

4:55-5:15 Sandwiching Random Regular Graphs Between Binomial Random Graphs

Brendan McKay, Australian National University, Australia; Mikhail Isaev, Monash University, Australia; *Pu Gao*, University of Waterloo, Canada

5:20-5:40 Factors and Loose Hamilton Cycles in Sparse Pseudo-Random Hypergraphs

Jie Han, University of Rhode Island, U.S.; Hiep Han, Universidad de Santiago de Chile, Chile; Patrick Morris, Freie Universitaet Berlin, Germany

5:45-6:05 Very Fast Construction of Bounded Degree Spanning Graphs Via the Semi-Random Graph Process

Omri Ben-Eliezer and Lior Gishboliner, Tel Aviv University, Israel; Dan Hefetz, Ariel University, Israel; Michael Krivelevich, Tel Aviv University, Israel

6:10-6:30 A New Algorithm for the Robust Semi-Random Independent Set Problem

Theo R. McKenzie and Hermish Mehta, University of California, Berkeley, U.S.; Luca Trevisan, Bocconi University, Italy

Intermission

6:35 p.m.-6:45 p.m.

ALENEX Business Meeting

6:45 p.m.-7:45 p.m.

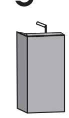

Room: Canyon A&B

Monday, January 6

Registration

8:00 a.m.-5:00 p.m.

Room: Topaz

Continental Breakfast

8:30 a.m.-9:00 a.m.

Room: Grand Ballroom Foyer

Monday, January 6

CP13

ALENEX Session 4

9:00 a.m.-11:05 a.m.

Room: Canyon A&B

Chair: Irene Finocchi, University of Rome La Sapienza, Italy

9:00-9:20 Engineering Top-Down Weight-Balanced Trees

Lukas Barth and Dorothea Wagner, Karlsruhe Institute of Technology, Germany

9:25-9:45 RecSplit: Minimal Perfect Hashing Via Recursive Splitting

Sebastiano Vigna, Universita degli Studi di Milano, Italy; Emmanuel Esposito, Università; degli Studi di Milano, Italy; *Thomas Mueller Graf*, Independent Scholar

9:50-10:10 Cost-Optimal Assignment of Elements in Genome-Scale Multi-Way Bucketed Cuckoo Hash Tables

Jens Zentgraf, Technical University of Dortmund, Germany; Henning Timm and Sven Rahmann, University of Duisburg-Essen, Germany

10:15-10:35 Reverse-Safe Data Structures for Text Indexing

Giulia Bernardini, University of Milano-Bicocca, Italy; Huiping Chen, King's College London, United Kingdom; Gabriele Fici, Universita degli Studi di Palermo - Italia; Grigorios Loukides, King's College London, United Kingdom; Solon P. Pissis, Centrum voor Wiskunde en Informatica (CWI), Netherlands

10:40-11:00 Constructing the Wavelet Tree and Wavelet Matrix in Distributed Memory

Patrick Dinklage, Technical University of Dortmund, Germany; Johannes Fischer and Florian Kurpicz, Technische Universität Dortmund, Germany

Monday, January 6

CP14

SODA Session 4A

9:00 a.m.-11:05 a.m.

Room: Grand Ballroom C

Chair: Sergio Cabello, University of Ljubljana, Slovenia

9:00-9:20 Tightening Curves on Surfaces Monotonically with Applications

Hsien-Chih Chang, Duke University, U.S.

9:25-9:45 Embeddability of Simplicial Complexes Is Undecidable

Marek Filakovsky, Institute of Science and Technology, Austria; Uli Wagner, Institute of Science and Technology Austria, Austria; Stephan Zhechev, Institute of Science and Technology, Austria

9:50-10:10 Optimal Bound on the Combinatorial Complexity of Approximating Polytopes

David M. Mount, University of Maryland, U.S.; Rahul Arya, University of California, Berkeley, U.S.; Sunil Arya, Hong Kong University of Science and Technology, Hong Kong; Guilherme da Fonseca, Aix-Marseille Université, France

10:15-10:35 Optimal Orthogonal Drawings of Planar 3-Graphs in Linear Time

Walter Didimo, *Giuseppe Liotta*, and Giacomo Ortali, Università di Perugia, Italy; Maurizio Patrignani, Universita di Roma Tre, Italy

10:40-11:00 Euclidean Bottleneck Bounded-Degree Spanning Tree Ratios

Ahmad Biniaz, University of Windsor, Canada

Monday, January 6

CP15

SODA Session 4B

9:00 a.m.-11:05 a.m.

Room: Alpine East

Chair: Karthekeyan Chandrasekara, University of Illinois at Urbana-Champaign, U.S.

9:00-9:20 Near-Optimal Approximate Discrete and Continuous Submodular Function Minimization

Yang Liu, Aaron Sidford, and *Brian Axelrod*, Stanford University, U.S.

9:25-9:45 Sticky Brownian Rounding and Its Applications to Constraint Satisfaction Problems

Sepehr Abbasi-Zadeh, University of Toronto, Canada; Nikhil Bansal, Centrum voor Wiskunde en Informatica (CWI), Netherlands; *Guru Guruganesh*, Google Research, U.S.; Aleksandar Nikolov, University of Toronto, Canada; Roy Schwartz, Technion, Israel; Mohit Singh, Georgia Institute of Technology, U.S.

9:50-10:10 Spherical Discrepancy Minimization and Algorithmic Lower Bounds for Covering the Sphere

Chris Jones and Matt McPartlon, University of Chicago, U.S.

10:15-10:35 Faster p-Norm Minimizing Flows, Via Smoothed q-Norm Problems

Deeksha Adil and *Sushant Sachdeva*, University of Toronto, Canada

10:40-11:00 Regular Languages Meet Prefix Sorting

Nicola Prezza, Universita di Pisa, Italy; Jarno Alanko, University of Helsinki, Finland; Giovanna D'Agostino and Alberto Policriti, Università di Udine, Italy

Monday, January 6

CP16

SODA Session 4C

9:00 a.m.-11:05 a.m.

Room: Alpine West

Chair: To Be Determined

9:00-9:20 Hitting Topological Minor Models in Planar Graphs is Fixed Parameter Tractable

Petr Golovach, University of Bergen, Norway; *Dimitrios Thilikos*, CNRS, LIRMM, France; Giannos Stamoulis, National and Kapodistrian University of Athens, Greece

9:25-9:45 A Complexity Dichotomy for Hitting Connected Minors on Bounded Treewidth Graphs: the Chair and the Banner Draw the Boundary

Julien Baste, Ulm University, Germany; Ignasi Sau and *Dimitrios Thilikos*, CNRS, LIRMM, France

9:50-10:10 Detecting Feedback Vertex Sets of Size k in O*ŕ(2.7^K) Time

Jason M. Li, Carnegie Mellon University, U.S.; Jesper Nederlof, Eindhoven University of Technology, Netherlands

10:15-10:35 A Nearly 5/3-Approximation Fpt Algorithm for Min-K-Cut

Ken-ichi Kawarabayashi, National Institute of Informatics, Japan; *Bingkai Lin*, Nanjing University, China

10:40-11:00 A $4 + \epsilon$ Approximation for k-Connected Subgraphs

Zeev Nutov, The Open University of Israel, Israel

Coffee Break

11:05 a.m.-11:30 a.m.

Room: Grand Ballroom Foyer

Monday, January 6

IP2

Dynamic Algorithms for Point Sets and Graphs

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

Room: Grand Ballroom C

Chair: Virginia Vassilevska Williams, Massachusetts Institute of Technology, U.S.

Many data sets change dynamically. If the data set is large, efficient algorithms are needed to handle the updates. If the data sets consists of elements with a binary relation, such as a social network graph, a dynamic graph algorithm can be used. If the data set consists of points in a metric space, such as high-dimensional feature vectors used for data mining, a dynamic data structure for point sets are needed. In this talk I will survey the state of the art of dynamic data structures for graphs and point sets. For both kinds of data sets, a hierarchical decomposition of the data set is often crucial for efficiency. Thus, I will present the decompositions used in the fastest known dynamic clustering algorithms for various definitions of clustering of point sets as well as two graph decompositions employed in dynamic graph algorithms for a large set of graph problems.

Monika Henzinger

University of Vienna, Austria

Lunch Break

12:30 p.m.-2:00 p.m.

Attendees on their own

Monday, January 6

CP17

SOSA Session 1

2:00 p.m.-4:05 p.m.

Room: Canyon A&B

Chair: Martin Farach-Colton, Rutgers University, U.S.

2:00-2:20 Reducing 3SUM to Convolution-3SUM

Timothy M. Chan and *Qizheng He*, University of Illinois at Urbana-Champaign, U.S.

2:25-2:45 Bucket Oblivious Sort: An Extremely Simple Oblivious Sort

Ling Ren, University of Illinois at Urbana-Champaign, U.S.

2:50-3:10 One (more) Line on the Most Ancient Algorithm in History

Ilya Volkovich, University of Michigan, U.S.; Bruno Grenet, Université de Montpellier and CNRS, France; *Vishwas Bhargava*, Rutgers University, U.S.

3:15-3:35 Multiplicative Rank-1 Approximation Using Length-Squared Sampling

Ragesh Jaiswal and Amit Kumar, Indian Institute of Technology, Delhi, India

3:40-4:00 Quantum Approximate Counting, Simplified

Scott Aaronson and *Patrick Rall*, University of Texas at Austin, U.S.

Monday, January 6

CP18

SODA Session 5A

2:00 p.m.-4:05 p.m.

Room: Grand Ballroom C

Chair: Karthekeyan Chandrasekara, University of Illinois at Urbana-Champaign, U.S.

2:00-2:20 2-Approximating Feedback Vertex Set in Tournaments

Saket Saurabh, Institute of Mathematical Sciences, India and University of Bergen, Norway; *Daniel Lokshtanov*, University of California, Santa Barbara, U.S.; Pranabendu Misra, Max Planck Institute for Informatics, Germany; Joydeep Mukherjee, Indian Statistical Institute, India; Fahad Panolan, Indian Institute of Technology Hyderabad, India; Geevarghese Philip, Chennai Mathematical Institute, India

2:25-2:45 Fast Lp-Based Approximations for Geometric Packing and Covering Problems

Kent Quanrud, Purdue University, U.S.; Chandra Chekuri and Sariel Har-Peled, University of Illinois at Urbana-Champaign, U.S.

2:50-3:10 Quasi-Polynomial Algorithms for Submodular Tree Orienteering and Other Directed Network Design Problems

Rohan Ghuge and Viswanath Nagarajan, University of Michigan, U.S.

3:15-3:35 Lossless Prioritized Embeddings

Michael Elkin, Ben-Gurion University, Israel

3:40-4:00 Labelings Vs. Embeddings: On Distributed Representations of Distances

Arnold Filtser, Columbia University, U.S.; Lee-Ad Gottlieb, Ariel University, Israel; Robert Krauthgamer, Weizmann Institute of Science, Israel

Monday, January 6

CP19

SODA Session 5B

2:00 p.m.-4:05 p.m.

Room: Alpine East

Chair: To Be Determined

2:00-2:20 Round Complexity of Common Randomness Generation: The Amortized Setting

Noah Golowich, Massachusetts Institute of Technology, U.S.; Madhu Sudan, Harvard University, U.S.

2:25-2:45 The Power of Distributed Verifiers in Interactive Proofs

Eylon Yogev, Simons Institute, U.S.; *Merav Parter* and Moni Naor, Weizmann Institute of Science, Israel

2:50-3:10 Lower Bounds for Oblivious Near-Neighbor Search

Kasper Green Larsen, Aarhus University, Denmark; Tal Malkin and Omri Weinstein, Columbia University, U.S.; *Kevin Yeo*, Google, Inc., U.S.

3:15-3:35 The Combinatorics of the Longest-Chain Rule: Linear Consistency for Proof-of-Stake Blockchains

Erica Blum, University of Maryland, College Park, U.S.; Aggelos Kiayias, IOHK and the University of Edinburgh, United Kingdom; Cristopher Moore, Santa Fe Institute, U.S.; *Saad Quader* and Alexander Russell, University of Connecticut, U.S.

3:40-4:00 A New Lower Bound on Hadwiger-Debrunner Numbers in the Plane

Chaya Keller, Ariel University, Israel; Shakhar Smorodinsky, Ben-Gurion University, Israel

Monday, January 6

CP20

SODA Session 5C

2:00 p.m.-4:05 p.m.

Room: Alpine West

Chair: To Be Determined

2:00-2:20 Navigating An Infinite Space with Unreliable Movements

Jara Uitto, Aalto University, Finland; *Anders Martinsson*, ETH Zürich, Switzerland

2:25-2:45 Linear Rankwidth Meets Stability

Patrice Ossona de Mendez, CNRS, France; Jaroslav Nesetril, Charles University, Czech Republic; Roman Rabinovich, Technische Universitaet Berlin, Germany; Sebastian Siebertz, University of Bremen, Germany

2:50-3:10 Edge Expansion and Spectral Gap of Nonnegative Matrices

Jenish C. Mehta and Leonard Schulman, California Institute of Technology, U.S.

3:15-3:35 Combinatorial Generation Via Permutation Languages

Elizabeth Hartung, Massachusetts College of Liberal Arts, U.S.; *Hung P. Hoang*, ETH Zürich, Switzerland; Torsten Mütze, University of Warwick, United Kingdom; Aaron Williams, Williams College, U.S.

3:40-4:00 X-Ramanujan Graphs

Sidhanth Mohanty, University of California, Berkeley, U.S.; Ryan O'Donnell, Carnegie Mellon University, U.S.

Coffee Break

4:05 p.m.-4:30 p.m.

Room: Grand Ballroom Foyer

Monday, January 6

CP21

SOSA Session 2

4:30 p.m.-6:35 p.m.

Room: Canyon A&B

Chair: Michael A. Bender, Stony Brook University, U.S.

4:30-4:50 Dynamic Generalized Closest Pair: Revisiting Eppstein's Technique

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

4:55-5:15 On the Change-Making Problem

Timothy M. Chan and *Qizheng He*, University of Illinois at Urbana-Champaign, U.S.

5:20-5:40 A Short Proof of the Toughness of Delaunay Triangulations

Ahmad Biniaz, University of Windsor, Canada

5:45-6:05 Adaptive Discrete Phase Retrieval

Moses Charikar, *Xian Wu*, and Yinyu Ye, Stanford University, U.S.

6:10-6:30 Fast Fourier Sparsity Testing

Samson Zhou, Carnegie Mellon University, U.S.; Grigory Yaroslavtsev, University of Pennsylvania, U.S.

Monday, January 6

CP22

SODA Session 6A

4:30 p.m.-6:35 p.m.

Room: Grand Ballroom C

Chair: Michael Dinitz, Johns Hopkins University, U.S.

4:30-4:50 Faster Deterministic Distributed Coloring Through Recursive List Coloring

Fabian Kuhn, Universität Freiburg, Germany

4:55-5:15 Faster Algorithms for Edge Connectivity Via Random 2-Out Contractions

Krzysztof Nowicki, University of Wroclaw, Poland; Mohsen Ghaffari, ETH Zürich, Switzerland; Mikkel Thorup, University of Copenhagen, Denmark

5:20-5:40 Shortest Paths in a Hybrid Network Model

John Augustine, Indian Institute of Technology Madras, India; Kristian Hinnenthal, Universität Paderborn, Germany; Fabian Kuhn, Universität Freiburg, Germany; Christian Scheideler, Universität Paderborn, Germany; *Philipp Schneider*, Universität Freiburg, Germany

5:45-6:05 Parallel Batch-Dynamic Graphs: Algorithms and Lower Bounds

Laxman Dhulipala, Carnegie Mellon University, U.S.; David Durfee, LinkedIn, U.S.; Janardhan Kulkarni, Microsoft Research, U.S.; Richard Peng and *Saurabh Sawlani*, Georgia Institute of Technology, U.S.; Xiaorui Sun, University of Ilinois at Chicago, U.S.

6:10-6:30 Finding a Bounded-Degree Expander Inside a Dense One

Francesco Pasquale, University of Rome II, Tor Vergata, Italy; Luca Becchetti, Universita di Roma "La Sapienza," Italy; Andrea Clementi, University of Rome II, Tor Vergata, Italy; Emanuele Natale, Université Côte d'Azur, INRIA, CNRS, France; Luca Trevisan, Bocconi University, Italy

Monday, January 6

CP23

SODA Session 6B

4:30 p.m.-6:35 p.m.

Room: Alpine East

Chair: To Be Determined

4:30-4:50 On Decoding Cohen-Haeupler-Schulman Tree Codes

Anand Kumar Narayanan, Sorbonne Universités, France; *Matthew Weidner*, Carnegie Mellon University, U.S.

4:55-5:15 On the Performance of Reed-Muller Codes with Respect to Random Errors and Erasures

Ori Sberlo and Amir Shpilka, Tel Aviv University, Israel

5:20-5:40 On the Power of Relaxed Local Decoding Algorithms

Tom Gur, University of Warwick, United Kingdom; *Oded Lachish*, University of London, United Kingdom

5:45-6:05 Relaxed Locally Correctable Codes with Nearly-Linear Block Length and Constant Query Complexity

Tom Gur, University of Warwick, United Kingdom

6:10-6:30 List Decoding of Direct Sum Codes

Vedat Levi Alev, University of Waterloo, Canada; *Fernando Granha Jeronimo* and Dylan Quintana, University of Chicago, U.S.; Shashank Srivastava and Madhur Tulsiani, Toyota Technological Institute at Chicago, U.S.

Monday, January 6

CP24

SODA Session 6C

4:30 p.m.-6:35 p.m.

Room: Alpine West

Chair: Julia Chuzhoy, Toyota Technological Institute at Chicago, U.S.

4:30-4:50 Improved Hardness for H-Colourings of G-Colourable Graphs

Marcin Wrochna and Stanislav Zivny, University of Oxford, United Kingdom

4:55-5:15 Ultimate Greedy Approximation of Independent Sets in Subcubic Graphs

Piotr J. Krysta, University of Liverpool, United Kingdom; Mathieu Mari, École Normale Supérieure Paris, France; Nan Zhi, University of Liverpool, United Kingdom

5:20-5:40 Counting Independent Sets in Unbalanced Bipartite Graphs

Sarah Cannon, Claremont McKenna College, U.S.; Will Perkins, University of Illinois at Chicago, U.S.

5:45-6:05 Faster Sublinear Approximations of k-Cliques for Low Arboricity Graphs

Talya Eden and Dana Ron, Tel Aviv University, Israel; C. Seshadhri, University of California, Santa Cruz, U.S.

6:10-6:30 Improved Inapproximability of Rainbow Coloring

Aditya Potukuchi, Rutgers University, U.S.; Per Austrin, KTH Royal Institute of Technology, Sweden; Amey Bhangale, Weizmann Institute of Science, Israel

Intermission

6:35 p.m.-6:45 p.m.

SODA Business Meeting & Awards Presentation, and SOSA Business Meeting

6:45 p.m.-7:45 p.m.

Room: Grand Ballroom C

Complimentary beer and wine will be served.

Tuesday, January 7

Registration

8:00 a.m.-5:00 p.m.

Room: Topaz

Continental Breakfast

8:30 a.m.-9:00 a.m.

Room: Grand Ballroom Foyer

Tuesday, January 7

CP25

SOSA Session 3

9:00 a.m.-11:05 a.m.

Room: Canyon A&B

Chair: Inge Li Gørtz, Technical University of Denmark, Denmark

9:00-9:20 Nearly Linear Time Approximations for Mixed Packing and Covering Problems Without Data Structures Or Randomization

Kent Quanrud, Purdue University, U.S.

9:25-9:45 Distributionally Robust Max Flows

Louis L. Chen, Naval Postgraduate School, U.S.; Will Ma, Columbia University, U.S.; James Orlin and David Simchi-Levi, Massachusetts Institute of Technology, U.S.

9:50-10:10 On a Decentralized (Δń+1) -Graph Coloring Algorithm

Paul B. De Supinski, Facebook, U.S.; Deeparnab Chakrabarty, Dartmouth College, U.S.

10:15-10:35 Distributed Backup Placement in One Round and its Applications to Maximum Matching Approximation and Self-Stabilization

Gal Oren, Ben Gurion University, Israel; Leonid Barenboim, Open University of Israel, Israel

10:40-11:00 Simple Label-Correcting Algorithms for Partially-Dynamic Approximate Shortest Paths in Directed Graphs

Adam Karczmarz, University of Warsaw, Poland; *Jakub Lacki*, Google Research, U.S.

Tuesday, January 7

CP26

SODA Session 7A

9:00 a.m.-11:05 a.m.

Room: Grand Ballroom C

Chair: Janardhan Kulkarni, Microsoft Research, U.S.

9:00-9:20 Chasing Nested Convex Bodies Nearly Optimally

Mark Sellke, Stanford University, U.S.; Yin Tat Lee, University of Washington, U.S.; Sébastien Bubueck, Microsoft Research, U.S.; Yuanzhi Li, Carnegie Mellon University, U.S.; Bo'az Klartag, Weizmann Institute of Science, Israel

9:25-9:45 Chasing Convex Bodies Optimally

Mark Sellke, Stanford University, U.S.

9:25-9:45 Chasing Convex Bodies with Linear Competitive Ratio

C.J. Argue and Anupam Gupta, Carnegie Mellon University, U.S.; Guru Guruganesh, Google Research, U.S.; Ziye Tang, Carnegie Mellon University, U.S.

9:50-10:10 The Online Submodular Cover Problem

Roie Levin and Anupam Gupta, Carnegie Mellon University, U.S.

10:15-10:35 Online Probabilistic Metric Embedding: A General Framework for Bypassing Inherent Bounds

Yair Bartal, Hebrew University, Israel; *Nova Fandina*, Hebrew University of Jerusalem, Israel; William Umboh, The University of Sydney, Australia

10:40-11:00 Achieving Optimal Backlog in the Vanilla Multi-Processor Cup Game

William Kuszmaul, Massachusetts Institute of Technology, U.S.

Tuesday, January 7

CP27

SODA Session 7B

9:00 a.m.-11:05 a.m.

Room: Alpine East

Chair: Virginia Vassilevska Williams, Massachusetts Institute of Technology, U.S.

9:00-9:20 Fine-Grained Complexity of Graph Homomorphism Problem for Bounded-Treewidth Graphs

Karolina Okrasa and *Pawel Rzazewski*, Warsaw University of Technology, Poland

9:25-9:45 Reducing Approximate Longest Common Subsequence to Approximate Edit Distance

Aviad Rubinstein, Harvard University, U.S.; *Zhao Song*, University of Texas at Austin, U.S.

9:50-10:10 Improved Algorithms for Edit Distance and LCS: Beyond Worst Case

Mahdi Boroujeni and Masoud Seddighin, Sharif University of Technology, Iran; *Saeed Seddighin*, Harvard University, U.S.

10:15-10:35 Hyperbolic Intersection Graphs and (quasi)-Polynomial Time

Sándor Kisfaludi-Bak, Max Planck Institute for Informatics, Germany

10:40-11:00 Adaptive Shivers Sort: An Alternative Sorting Algorithm

Vincent Jugé, Université Paris-Est Marne-la-Vallée, France

Tuesday, January 7

CP28

SODA Session 7C

9:00 a.m.-11:05 a.m.

Room: Alpine West

Chair: Sepehr Assadi, Princeton University and Rutgers University, U.S.

9:00-9:20 Tight Bounds for the Subspace Sketch Problem with Applications

Yi Li, Nanyang Technological University, Singapore; Ruosong Wang and *David Woodruff*, Carnegie Mellon University, U.S.

9:25-9:45 Composable Core-Sets for Determinant Maximization Problems Via Spectral Spanners

Piotr Indyk, Massachusetts Institute of Technology, U.S.; Sepideh Mahabadi, Toyota Technological Institute at Chicago, U.S.; Shayan Oveis Gharan and *Alireza Rezaei*, University of Washington, U.S.

9:50-10:10 New (α, β) Spanners and Hopsets

Merav Parter and Uri Ben-Levy, Weizmann Institute of Science, Israel

10:15-10:35 The Communication Complexity of Set Intersection and Multiple Equality Testing

Dawei Huang and Seth Pettie, University of Michigan, U.S.; Zhijun Zhang and Yixiang Zhang, Tsinghua University, P. R. China

10:40-11:00 The Communication Complexity of Optimization

Santosh Vempala, Georgia Institute of Technology, U.S.; Ruosong Wang and *David Woodruff*, Carnegie Mellon University, U.S.

Coffee Break

11:05 a.m.-11:30 a.m.

Room: Grand Ballroom Foyer

Tuesday, January 7

IP3

Computing Maximum Independent Set in Graph Classes

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

Room: Grand Ballroom C

Chair: Zdenek Dvorak, Charles University, Czech Republic

The Maximum Independent Set (MIS) in graphs is among the most classical algorithmically hard problems. However, it can sometimes be convenient to code a question in terms of MIS and try to take advantage of the graph structure. As an example, the only known method to approximate (PTAS) a maximum collection of pairwise intersecting disks in the plane is to form the intersection graph and solve MIS on its complement. The crucial point is that such a MIS instance belongs to a particular graph class, and indeed the algorithm is heavily guided by the forbidden induced subgraphs to discover a good approximation of MIS. Oddly, the geometric structure is completely ignored in the process. In this talk I will mainly focus on the approximation of MIS in graph classes defined by a finite list of forbidden subgraphs. A particular emphasis will be made on the interplay between graph structure and algorithms, and how classical conjectures (like for instance Erdos-Hajnal) fit in this landscape. I will also discuss some (easy) tools which are surprisingly efficient as long as one is happy with quasi-polynomial time approximation. Polytime algorithms in this area remain however a complete mystery.

Stéphan Thomassé

École Normale Supérieure de Lyon, France

Lunch Break

12:30 p.m.-2:00 p.m.

Attendees on their own

Tuesday, January 7

CP29

SODA Session 8A

2:00 p.m.-4:05 p.m.

Room: Grand Ballroom C

Chair: Sepehr Assadi, Princeton University and Rutgers University, U.S.

2:00-2:20 Space Efficient Approximation to Maximum Matching Size from Uniform Edge Samples

Jakab Tardos, EPFL, Switzerland; Michael Kapralov, École Polytechnique Fédérale de Lausanne, Switzerland; Slobodan Mitrovic, Massachusetts Institute of Technology, U.S.; Ashkan Norouzi-Fard, Google Zurich, Switzerland

2:25-2:45 Approximate Maximum Matching in Random Streams

Alireza Farhadi, University of Maryland, U.S.; MohammadTaghi Hajiaghayi, University of Maryland, College Park, U.S.; Tung Mai, Anup Rao, and Ryan Rossi, Adobe Systems, U.S.

2:50-3:10 Vertex Ordering Problems in Directed Graph Streams

Amit Chakrabarti and *Prantar Ghosh*, Dartmouth College, U.S.; Andrew McGregor, University of Massachusetts, Amherst, U.S.; Sofya Vorotnikova, Dartmouth College, U.S.

3:15-3:35 Faster Update Time for Turnstile Streaming Algorithms

Josh Alman, Harvard University, U.S.; Huacheng Yu, Princeton University, U.S.

3:40-4:00 Fast and Space Efficient Spectral Sparsification in Dynamic Streams

Jakab Tardos, EPFL, Switzerland; Michael Kapralov, École Polytechnique Fédérale de Lausanne, Switzerland; Aida Mousavifar, EPFL, Switzerland; Cameron Musco, University of Massachusetts, Amherst, U.S.; Christopher Musco, New York University, U.S.; Navid Nouri, EPFL, Switzerland; Aaron Sidford, Stanford University, U.S.

Tuesday, January 7

CP30

SODA Session 8B

2:00 p.m.-4:05 p.m.

Room: Alpine East

Chair: Sungjin Im, University of California, Merced, U.S.

2:00-2:20 Near-Optimal Bounds for Online Caching with Machine Learned Advice

Dhruv W. Rohatgi, Massachusetts Institute of Technology, U.S.

2:25-2:45 Interleaved Caching with Access Graphs

Ravi Kumar, Google, Inc., U.S.; *Manish Purohit*, Zoya Svitkina, and Erik Vee, Google Research, U.S.

2:50-3:10 Online Scheduling Via Learned Weights

Silvio Lattanzi, Google Zurich, Switzerland; *Thomas J. Lavastida* and Benjamin Moseley, Carnegie Mellon University, U.S.; Sergei Vassilvitskii, Google Research, U.S.

3:15-3:35 Competitive Online Search Trees on Trees

Grigorios Koumoutsos, Université Libre de Bruxelles, Belgium; Prosenjit Bose, Carleton University, Canada; Jean Cardinal, Université Libre de Bruxelles, Belgium; John Iacono, Polytechnic Institute of New York University, U.S.; Stefan Langerman, Université Libre de Bruxelles, Belgium

3:40-4:00 Quantifying the Burden of Exploration and the Unfairness of Free Riding

Christopher Jung and Sampath Kannan, University of Pennsylvania, U.S.; *Neil Lutz*, Iowa State University, U.S.

Tuesday, January 7

CP31

SODA Session 8C

2:00 p.m.-4:05 p.m.

Room: Alpine West

Chair: Julia Chuzhoy, Toyota Technological Institute at Chicago, U.S.

2:00-2:20 Diameter Computation on H-Minor Free Graphs and Graphs of Bounded (distance) Vc-Dimension

Guillaume Ducoffe, ICI – National Institute for Research and Development informatics, Romania; Michel Habib, Université Paris, France; Laurent Viennot, Inria, France

2:25-2:45 A Strongly Polynomial Algorithm for Finding a Shortest Non-Zero Path in Group-Labeled Graphs

Yutaro Yamaguchi, Osaka University, Japan

2:50-3:10 A Blossom Algorithm for Maximum Edge-Disjoint T-Paths

Satoru Iwata, University of Tokyo, Japan; Yu Yokoi, National Institute of Informatics, Japan

3:15-3:35 A Face Cover Perspective to ℓ_1 Embeddings of Planar Graphs

Arnold Filtser, Columbia University, U.S.

3:40-4:00 Multi-Transversals for Triangles and the Tuza Conjecture

Parinya Chalermsook, Aalto University, Finland; Samir Khuller and Pattara Sukprasert, Northwestern University, U.S.; *Sumedha Uniyal*, Aalto University, Finland

TUNGA Workshop

(see *http://tunga.eu/* for details)

2:00 p.m.-4:05 p.m.

Room: Canyon A&B

Coffee Break

4:05 p.m.-4:30 p.m.

Room: Grand Ballroom Foyer

Tuesday, January 7

CP32

SODA Session 9A

4:30 p.m.-6:35 p.m.

Room: Grand Ballroom C

Chair: Michael Kapralov, École Polytechnique Fédérale de Lausanne, Switzerland

4:30-4:50 Domain Reduction for Monotonicity Testing: A $o(d)$ Tester for Boolean Functions in d-Dimensions

Hadley Black, University of California, Los Angeles, U.S.; Deeparnab Chakrabarty, Dartmouth College, U.S.; C. Seshadhri, University of California, Santa Cruz, U.S.

4:55-5:15 Approximating the Distance to Monotonicity of Boolean Functions

Erik Waingarten, Columbia University, U.S.; Ramesh Krishnan S. Pallavoor and Sofya Raskhodnikova, Boston University, U.S.

5:20-5:40 Reconstruction under Outliers for Fourier-Sparse Functions

Xue Chen, Northwestern University, U.S.; Anindya De, University of Pennsylvania, U.S.

5:45-6:05 Testing Convexity of Functions over Finite Domains

Eric Blais, University of Waterloo, Canada; Aleksandrs Belovs, University of Latvia, Latvia; Abhinav Bommireddi, University of Waterloo, Canada

6:10-6:30 Computing and Testing Small Connectivity in Near-Linear Time and Queries Via Fast Local Cut Algorithms

Sebastian Forster, University of Salzburg, Germany; Danupon Nanongkai, KTH Royal Institute of Technology, Sweden; Thatchaphol Saranurak, Toyota Technological Institute at Chicago, U.S.; Liu Yang, Independent Researcher; *Sorrachai Yingchareonthawornchai*, Aalto University, Finland

Tuesday, January 7

CP33

SODA Session 9B

4:30 p.m.-6:35 p.m.

Room: Alpine East

Chair: Michael Dinitz, Johns Hopkins University, U.S.

4:30-4:50 The Two-Sided Game of Googol and Sample-Based Prophet Inequalities

José Correa, Andrés Cristi, *Boris Epstein*, and José Soto, Universidad de Chile, Chile

4:55-5:15 Competitive Analysis with a Sample and the Secretary Problem

Haim Kaplan, Tel Aviv University, Israel; *David Naori*, Technion Israel Institute of Technology, Israel; Danny Raz, Bar-Ilan University, Israel

5:20-5:40 A Truthful Cardinal Mechanism for One-Sided Matching

Vasilis Gkatzelis, Drexel University, U.S.; Rediet Abebe, Cornell University and Harvard University, U.S.; Richard Cole, New York University, U.S.; *Jason Hartline*, Northwestern University, U.S.

5:45-6:05 Cake Cutting on Graphs: A Discrete and Bounded Proportional Protocol

Zhijie Zhang, Chinese Academy of Sciences, China; Xiaohui Bei, Nanyang Technological University, Singapore; Xiaoming Sun, Hao Wu, Jialin Zhang, and Wei Zi, Chinese Academy of Sciences, China

6:10-6:30 Instance-Optimality in the Noisy Value-and Comparison-Model

Frederik Mallmann-Trenn, King's College London, United Kingdom; Claire Mathieu, CNRS, France; Vincent Cohen-Addad, Sorbonne Universités and CNRS, France

Tuesday, January 7

CP34

SODA Session 9C

4:30 p.m.-6:35 p.m.

Room: Alpine West

Chair: David Woodruff, Carnegie Mellon University, U.S.

4:30-4:50 Reconstruction of Depth-4 Multilinear Circuits

Vishwas Bhargava and Shubhangi Saraf, Rutgers University, U.S.; Ilya Volkovich, University of Michigan, U.S.

4:55-5:15 Counting and Finding Homomorphisms Is Universal for Parameterized Complexity Theory

Marc Roth, Oxford University, United Kingdom; *Philip Wellnitz*, Max Planck Institute for Informatics, Germany

5:20-5:40 Parameterized Complexity and Approximability of Directed Odd Cycle Transversal

Saket Saurabh, Institute of Mathematical Sciences, India and University of Bergen, Norway; *Daniel Lokshtanov*, University of California, Santa Barbara, U.S.; Meirav Zehavi, Ben Gurion University, Israel; Ramanujan M. S., University of Warwick, United Kingdom

5:45-6:05 Approximately Counting and Sampling Small Witnesses Using a Colourful Decision Oracle

Holger Dell, IT University of Copenhagen, Denmark; *John Lapinskas*, University of Bristol, United Kingdom; Kitty Meeks, University of Glasgow, Scotland, United Kingdom

6:10-6:30 Improved Bounds for Centered Colorings

Michal Debski, Warsaw University of Technology, Poland and Masaryk University, Brno, Czech Republic; Stefan Felsner, Technische Universität, Berlin, Germany; Piotr Micek, Jagiellonian University, Poland; *Felix Schröder*, Technische Universitaet Berlin, Germany

Tuesday, January 7

TUNGA Workshop

(see *http://tunga.eu/* for details)

4:30 p.m.-6:35 p.m.

Room: Canyon A&B

Intermission

6:35 p.m.-6:45 p.m.

Tuesday, January 7

IP4

SIAG/ACDA Plenary Presentation - Models and Algorithms in Support of Modern Urban Mobility

6:45 p.m.-7:45 p.m.

Room: Grand Ballroom C

Chair: Alex Pothen, Purdue University, U.S.

The sharing economy is transforming how people move within cities. Ride-sharing, bike-sharing (both station-based and dockless, with both traditional and electric pedal-assist bikes), and scooters have augmented the options present for decades in terms of public transit and private vehicle ownership. We will discuss the issues raised by this transformation, while focusing on some of the operational challenges present in station-based bike-sharing systems. We have worked with Citibike (the operator of bike-sharing in NYC) and its then parent company Motivate, developing optimization models and algorithms to change how they manage their systems. For example, continuous-time Markov chain models, combined with simple mathematical programming tools, can be used to answer the question – what is the optimal deployment of the bike fleet across the system at the start of the day? Furthermore, we consider the more strategic question of how to (re-) allocate dock-capacity in such systems, which gives rise to new (solvable) integer programming problems. We have also guided the development of Bike Angels, a program to incentivize users to crowdsource "rebalancing rides"; we will describe its underlying analytics, where the pricing mechanism is once again grounded in the same underlying algorithmic tools. This is joint work primarily with Daniel Freund, Shane Henderson, and Eoin O'Mahony.

David B. Shmoys
Cornell University, U.S.

Wednesday, January 8

Registration

8:00 a.m.-5:00 p.m.

Room: Topaz

Continental Breakfast

8:30 a.m.-9:00 a.m.

Room: Grand Ballroom Foyer

Wednesday, January 8

CP35

APOCS Session 1

9:00 a.m.-11:05 a.m.

Room: Canyon A&B

Chair: Yan Gu, University of California, Riverside, U.S.

9:00-9:20 Writeback-Aware Caching

Charles Mcguffey, Nathan Beckmann, Phillip Gibbons, and Bernhard Haeupler, Carnegie Mellon University, U.S.

9:25-9:45 Parallel Algorithms for Butterfly Computations

Jessica Shi and Julian Shun, Massachusetts Institute of Technology, U.S.

9:50-10:10 Memory-Efficient Performance Monitoring on Programmable Switches with Lean Algorithms

Zaoxing Liu and Samson Zhou, Carnegie Mellon University, U.S.; Ori Rottenstreich, Technion, Israel; Vladimir Braverman, Johns Hopkins University, U.S.; Jennifer Rexford, Princeton University, U.S.

10:15-10:35 Lock-Free Hopscotch Hashing

Robert Kelly, Barak Pearlmutter, and Phil Maguire, National University of Ireland, Maynooth, Ireland

10:40-11:00 Learning Software Constraints Via Installation Attempts

Ran Ben Basat, Harvard University, U.S.; Maayan Goldstein and Itai Segall, Nokia Bell Labs, U.S.

Wednesday, January 8

CP36

SODA Session 10A

9:00 a.m.-11:05 a.m.

Room: Grand Ballroom C

Chair: Sungjin Im, University of California, Merced, U.S.

9:00-9:20 Baker Game and Polynomial-Time Approximation Schemes

Zdenek Dvorak, Charles University, Czech Republic

9:25-9:45 Approximation Schemes for Capacitated Clustering in Doubling Metrics

Vincent Cohen-Addad, Sorbonne Universités and CNRS, France

9:50-10:10 Quasi-Polynomial Time Approximation Schemes for the Maximum Weight Independent Set Problem in *H*-Free Graphs

Marcin Pilipczuk, University of Warsaw, Poland; Maria Chudnovsky, Princeton University, U.S.; Michal Pilipczuk, University of Warsaw, Poland; *Stéphan Thomassé*, École Normale Supérieure de Lyon, France

10:15-10:35 A Ptas for Subset Tsp in Minor-Free Graphs

Hung Le, University of Victoria, Canada

10:40-11:00 Approximation Schemes Via Width/Weight Trade-Offs on Minor-Free Graphs

Saket Saurabh, Institute of Mathematical Sciences, India and University of Bergen, Norway; Fedor V. Fomin, University of Bergen, Norway; *Daniel Lokshtanov*, University of California, Santa Barbara, U.S.; Meirav Zehavi, Ben Gurion University, Israel

Wednesday, January 8

CP37

SODA Session 10B

9:00 a.m.-11:05 a.m.

Room: Alpine East

Chair: Artur Czumaj, University of Warwick, United Kingdom

9:00-9:20 Efficiently List-Edge Coloring Multigraphs Asymptotically Optimally

Fotis Iliopoulos and Alistair Sinclair, University of California, Berkeley, U.S.

9:25-9:45 Linear Size Sparsifier and the Geometry of the Operator Norm Ball

Victor Reis and Thomas Rothvoss, University of Washington, U.S.

9:50-10:10 Small Memory Robust Simulation of Client-Server Interactive Protocols over Oblivious Noisy Channels

T-H. Hubert Chan, University of Hong Kong, Hong Kong; Zhibin Liang, University of Hong Kong, China; *Antigoni Polychroniadou*, J.P. Morgan, U.S.; Elaine Shi, Cornell University, U.S.

10:15-10:35 Finding Perfect Matchings in Dense Hypergraphs

Jie Han, University of Rhode Island, U.S.; Peter Keevash, University of Oxford, United Kingdom

10:40-11:00 Worst-Case Polylog Incremental Spqr-Trees: Embeddings, Planarity, and Triconnectivity

Eva Rotenberg, Technical University of Denmark, Denmark; *Jacob Holm*, University of Copenhagen, Denmark

Wednesday, January 8

CP38

SODA Session 10C

9:00 a.m.-11:05 a.m.

Room: Alpine West

Chair: To Be Determined

9:00-9:20 Dominantly Truthful Multi-Task Peer Prediction, with Constant Number of Tasks

Yuqing Kong, Peking University, China

9:25-9:45 Selling Information Through Consulting

Yiling Chen, Harvard University, U.S.; Haifeng Xu, University of Virginia, U.S.; *Shuran Zheng*, Harvard University, U.S.

9:50-10:10 Algorithmic Price Discrimination

Rachel Cummings, Georgia Institute of Technology, U.S.; Nikhil R. Devanur, Microsoft Research, U.S.; Zhiyi Huang and *Xiangning Wang*, University of Hong Kong, Hong Kong

10:15-10:35 Bulow-Klemperer-Style Results for Welfare Maximization in Two-Sided Markets

Moshe Babaioff, Microsoft Research, U.S.; Kira Goldner, Columbia University, U.S.; *Yannai A. Gonczarowski*, Microsoft Research, U.S.

10:40-11:00 Inference from Auction Prices

Jason Hartline, Northwestern University, U.S.; *Aleck Johnsen*, Northwestern University, U.S.; Denis Nekipelov, University of Virginia, U.S.; Zihe Wang, Shanghai University of Finance and Economics, China

Coffee Break

11:05 a.m.-11:30 a.m.

Room: Grand Ballroom Foyer

Wednesday, January 8

IP5

Phase Transitions in Random Constraint Satisfaction Problems

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

Room: Grand Ballroom C

Chair: Shuchi Chawla, University of Wisconsin, Madison, U.S.

I will survey recent progress in determination of asymptotic behavior for random constraint satisfaction problems. This includes phase transitions and some understanding of solution geometry. This lecture is based in part on joint works with Zsolt Bartha, Jian Ding, Allan Sly, and Yumeng Zhang; I will also survey related results in the area.

Nike Sun
Massachusetts Institute of Technology, U.S.

Lunch Break

12:30 p.m.-2:00 p.m.

Attendees on their own

Wednesday, January 8

CP39

APOCS Session 2

2:00 p.m.-4:05 p.m.

Room: Canyon A&B

Chair: Tao Schardl, Massachusetts Institute of Technology, U.S.

2:00-2:20 Eccentricity Heuristics Through Sublinear Analysis Lenses

Tal Wagner, Massachusetts Institute of Technology, U.S.

2:25-2:45 Fast Distributed Backup Placement in Sparse and Dense Networks

Gal Oren, Ben Gurion University, Israel; Leonid Barenboim, Open University of Israel, Israel

2:50-3:10 Improved Parallel Cache-Oblivious Algorithms for Dynamic Programming

Yan Gu, University of California, Riverside, U.S.; Guy Blelloch, Carnegie Mellon University, U.S.

3:15-3:35 Online Flow Computation on Unit-Vertex-Capacitated Networks

Makis Arsenis and Robert Kleinberg, Cornell University, U.S.

3:40-4:00 Two-Way Trees: A Distributed Algorithm for Efficient Replica Search and Placement

Gahyun Park, Minseok Kwon, and Ramprasad Tamilselvan, Rochester Institute of Technology, U.S.; Seungjoon Lee, Google, Inc., U.S.

Wednesday, January 8

CP40

SODA Session 11A

2:00 p.m.-4:05 p.m.

Room: Grand Ballroom C

Chair: Merav Parter, Weizmann Institute of Science, Israel

2:00-2:20 Fully Dynamic Matching: Beating 2-Approximation in Δ^{ϵ} Update Time

Soheil Behnezhad, University of Maryland, U.S.; Jakub Lacki, Google Research, U.S.; Vahab Mirrokni, Google, Inc., U.S.

2:25-2:45 An Improved Algorithm for Incremental Cycle Detection and Topological Ordering in Sparse Graphs

Sayan Bhattacharya, University of Warwick, United Kingdom; Janardhan Kulkarni, Microsoft Research, U.S.

2:50-3:10 Deterministic Algorithms for Decremental Approximate Shortest Paths: Faster and Simpler

Maximilian Probst Gutenberg and Christian Wulff-Nilsen, University of Copenhagen, Denmark

3:15-3:35 Decremental SSSP in Weighted Digraphs: Faster and Against an Adaptive Adversary

Maximilian Probst Gutenberg and *Christian Wulff-Nilsen*, University of Copenhagen, Denmark

3:40-4:00 Fully-Dynamic All-Pairs Shortest Paths: Improved Worst-Case Time and Space Bounds

Maximilian Probst Gutenberg and Christian Wulff-Nilsen, University of Copenhagen, Denmark

Wednesday, January 8

CP41

SODA Session 11B

2:00 p.m.-4:05 p.m.

Room: Alpine East

Chair: To Be Determined

2:00-2:20 On the Cover of the Rolling Stone

Adrian Dumitrescu, University of Wisconsin, Milwaukee, U.S.; *Csaba Toth*, California State University, Northridge, U.S.

2:25-2:45 Computing Minimal Persistent Cycles: Polynomial and Hard Cases

Tamal K. Dey, *Tao Hou*, and Sayan Mandal, Ohio State University, U.S.

2:50-3:10 The Impacts of Dimensionality, Diffusion, and Directedness on Intrinsic Universality in the Abstract Tile Assembly Model

Daniel Hader, Aaron Koch, Matthew Patitz, and Michael Sharp, University of Arkansas, U.S.

3:15-3:35 Hierarchical Shape Construction and Complexity for Slidable Polyominos under Uniform External Forces

Tim Wylie, Jose Balanza-Martinez, David Caballero, Angel Cantu, Mauricio Flores, *Timothy Gomez*, Austin Luchsinger, Rene Reyes, and Robert Schweller, University of Texas, Rio Grande Valley

3:40-4:00 Even Maps, the Colin de Verdière Number and Representations of Graphs

Vojtech Kaluža, Universität Innsbruck, Austria; Martin Tancer, Charles University in Prague, Czech Republic

Wednesday, January 8

CP42

SODA Session 11C

2:00 p.m.-4:05 p.m.

Room: Alpine West

Chair: Shuchi Chawla, University of Wisconsin, Madison, U.S.

2:00-2:20 A Little Charity Guarantees Almost Envy-Freeness

Bhaskar R. Chaudhury, Max Planck Institute for Informatics, Germany; Tellikepalli Kavitha, Tata Institute of Fundamental Research, India; Kurt Mehlhorn, Max Planck Institute for Informatics, Germany; Alkmini Sgouritsa, University of Liverpool, United Kingdom

2:25-2:45 Approximating Nash Social Welfare under Submodular Valuations Through (Un)Matchings

Rucha Kulkarni, Jugal Garg, and Pooja Kulkarni, University of Illinois at Urbana-Champaign, U.S.

2:50-3:10 The Complexity of Contracts

Paul Dütting, London School of Economics, United Kingdom; Tim Roughgarden, Columbia University, U.S.; Inbal Talgam-Cohen, Technion, Israel

3:15-3:35 On the Tractability of Public Persuasion with No Externalities

Haifeng Xu, University of Virginia, U.S.

3:40-4:00 Complexity and Parametric Computation of Equilibria in Atomic Splittable Congestion Games Via Weighted Block Laplacians

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

Coffee Break

4:05 p.m.-4:30 p.m.

Room: Grand Ballroom Foyer

Wednesday, January 8

CP43

APOCS Session 3, followed by Business Meeting

4:30 p.m.-6:35 p.m.

Room: Canyon A&B

Chair: Phillip Gibbons, Carnegie Mellon University, U.S.

4:30-4:50 Scheduling I/O Latency-Hiding Futures in Task-Parallel Platforms

Kyle Singer, Kunal Agrawal, and I-Ting Angelina Lee, Washington University in St. Louis, U.S.

4:55-5:15 Cilkmem: Algorithms for Analyzing the Memory High-Water Mark of Fork-Join Parallel Programs

Tim Kaler, *William Kuszmaul*, Tao Schardl, and Daniele Vettorel, Massachusetts Institute of Technology, U.S.

5:20 APOCS Business Meeting

Wednesday, January 8

CP44

SODA Session 12A

4:30 p.m.-6:35 p.m.

Room: Grand Ballroom C

Chair: Sungjin Im, University of California, Merced, U.S.

4:30-4:50 A Tale of Santa Claus, Hypergraphs and Matroids

Sami Davies, Thomas Rothvoss, and Yihao Zhang, University of Washington, U.S.

4:55-5:15 Parallel Machine Scheduling to Minimize Energy Consumption

Nikhil Kumar, Indian Institute of Technology, Delhi, India; Antonios Antoniadis, Max Planck Institute for Informatics, Germany; Naveen Garg, Indian Institute of Technology, Delhi, India; Gunjan Kumar, Tata Institute of Fundamental Research, India

5:20-5:40 Hierarchy-Based Algorithms for Minimizing Makespan under Precedence and Communication Constraints

Janardhan Kulkarni, Microsoft Research, U.S.; Shi Li, State University of New York at Buffalo, U.S.; *Jakub Tarnawski*, Microsoft Research, U.S.; Minwei Ye, University at Buffalo, U.S.

5:45-6:05 Weighted Completion Time Minimization for Unrelated Machines Via Iterative Fair Contention Resolution

Sungjin Im and Maryam Shadloo, University of California, Merced, U.S.

6:10-6:30 How to Aggregate Top-Lists: Approximation Algorithms Via Scores and Average Ranks

Simon Mauras, Université de Paris, IRIF, CNRS, France; Claire Mathieu, CNRS, France

Wednesday, January 8

CP45

SODA Session 12B

4:30 p.m.-6:35 p.m.

Room: Alpine East

Chair: To Be Determined

4:30-4:50 Connectivity of Triangulation Flip Graphs in the Plane (Part I: Edge Flips)

Uli Wagner, Institute of Science and Technology Austria, Austria; Emo Welzl, ETH Zürich, Switzerland

4:55-5:15 Nearly Optimal Planar k Nearest Neighbors Queries under General Distance Functions

Chih-Hung Liu, ETH Zürich, Switzerland

5:20-5:40 Computing Circle Packing Representations of Planar Graphs

Sally Dong and Yin Tat Lee, University of Washington, U.S.; Kent Quanrud, University of Illinois at Urbana-Champaign, U.S.

5:45-6:05 Atomic Embeddability, Clustered Planarity, and Thickenability

Radoslav Fulek, University of Arizona, U.S.

6:10-6:30 The Stable Set Problem in Graphs with Bounded Genus and Bounded Odd Cycle Packing Number

Stefan Weltge, Technical University of Munich, Germany; Michele Conforti, Università degli Studi di Padova, Italy; Samuel Fiorini, Tony Huynh, and Gwenaël Joret, Université Libre de Bruxelles, Belgium

Wednesday, January 8

CP46

SODA Session 12C

4:30 p.m.-6:35 p.m.

Room: Alpine West

Chair: Michael Kapralov, École Polytechnique Fédérale de Lausanne, Switzerland

4:30-4:50 Nearly Optimal Edge Estimation with Independent Set Queries

Amit Levi, University of Waterloo, Canada; Xi Chen and Erik Waingarten, Columbia University, U.S.

4:55-5:15 A Lower Bound on Cycle-Finding in Sparse Digraphs

Timothy Sun, Xi Chen, Tim Randolph, and Rocco A. Servedio, Columbia University, U.S.

5:20-5:40 Robust Clustering Oracle and Local Reconstructor of Cluster Structure of Graphs

Pan Peng, University of Sheffield, United Kingdom

5:45-6:05 Sublinear Time Approximation of the Cost of a Metric k-Nearest Neighbor Graph

Artur Czumaj, University of Warwick, United Kingdom; Christian Sohler, University of Cologne, Germany

6:10-6:30 Improved Local Computation Algorithm for Set Cover Via Sparsification

Christoph Grunau, ETH Zürich, Switzerland; Slobodan Mitrovic, Massachusetts Institute of Technology, U.S.; Ronitt Rubinfeld, Massachusetts Institute of Technology, U.S. and Tel Aviv University, Israel; Ali Vakilian, Massachusetts Institute of Technology, U.S.

Abstracts

SIAM 2020

ACM-SIAM Symposium on
Discrete Algorithms
January 5–8, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Algorithm Engineering and Experiments
January 5–6, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Simplicity in Algorithms
January 6–7, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Algorithmic Principles of Computer Systems
January 8, 2020 • Salt Lake City, Utah, U.S.

IP1

Algorithmic Questions in Machine Learning

Machine learning (ML) is a very algorithmic discipline. But to experts in traditional algorithms (e.g. from STOC/FOCS/SODA) it is not always clear how to work on ML problems —especially deep learning— due to a sizeable gulf in mindsets as well as well-known challenges in thinking about nonconvex optimization. This talk will use a series of vignettes to illustrate interesting theory questions in contemporary ML, and how algorithmic insights have been obtained in recent years.

Sanjeev Arora
Princeton University
Computer Science Dept.
arora@cs.princeton.edu

IP2

Dynamic Algorithms for Point Sets and Graphs

Many data sets change dynamically. If the data set is large, efficient algorithms are needed to handle the updates. If the data sets consists of elements with a binary relation, such as a social network graph, a dynamic graph algorithm can be used. If the data set consists of points in a metric space, such as high-dimensional feature vectors used for data mining, a dynamic data structure for point sets are needed. In this talk I will survey the state of the art of dynamic data structures for graphs and point sets. For both kinds of data sets, a hierarchical decomposition of the data set is often crucial for efficiency. Thus, I will present the decompositions used in the fastest known dynamic clustering algorithms for various definitions of clustering of point sets as well as two graph decompositions employed in dynamic graph algorithms for a large set of graph problems.

Monika Henzinger
University of Vienna, Austria
monika.henzinger@univie.ac.at

IP3

Computing Maximum Independent Set in Graph Classes

he Maximum Independent Set (MIS) in graphs is among the most classical algorithmically hard problems. However, it can sometimes be convenient to code a question in terms of MIS and try to take advantage of the graph structure. As an example, the only known method to approximate (PTAS) a maximum collection of pairwise intersecting disks in the plane is to form the intersection graph and solve MIS on its complement. The crucial point is that such a MIS instance belongs to a particular graph class, and indeed the algorithm is heavily guided by the forbidden induced subgraphs to discover a good approximation of MIS. Oddly, the geometric structure is completely ignored in the process. In this talk I will mainly focus on the approximation of MIS in graph classes defined by a finite list of forbidden subgraphs. A particular emphasis will be made on the interplay between graph structure and algorithms, and how classical conjectures (like for instance Erdos-Hajnal) fit in this landscape. I will also discuss some (easy) tools which are surprisingly efficient as long as one is happy with quasi-polynomial time approximation. Polytime algorithms in this area remain however a complete mystery.

Stéphan Thomassé
ENS de Lyon
stephan.thomasse@ens-lyon.fr

IP4

SIAG/ACDA Plenary Presentation - Models and Algorithms in Support of Modern Urban Mobility

The sharing economy is transforming how people move within cities. Ride-sharing, bike-sharing (both station-based and dockless, with both traditional and electric pedal-assist bikes), and scooters have augmented the options present for decades in terms of public transit and private vehicle ownership. We will discuss the issues raised by this transformation, while focusing on some of the operational challenges present in station-based bike-sharing systems. We have worked with Citibike (the operator of bike-sharing in NYC) and its then parent company Motivate, developing optimization models and algorithms to change how they manage their systems. For example, continuous-time Markov chain models, combined with simple mathematical programming tools, can be used to answer the question what is the optimal deployment of the bike fleet across the system at the start of the day? Furthermore, we consider the more strategic question of how to (re-)allocate dock-capacity in such systems, which gives rise to new (solvable) integer programming problems. We have also guided the development of Bike Angels, a program to incentivize users to crowdsource rebalancing rides; we will describe its underlying analytics, where the pricing mechanism is once again grounded in the same underlying algorithmic tools. This is joint work primarily with Daniel Freund, Shane Henderson, and Eoin OMahony.

David B. Shmoys
Cornell University
david.shmoys@cornell.edu

IP5

Phase Transitions in Random Constraint Satisfaction Problems

I will survey recent progress in determination of asymptotic behavior for random constraint satisfaction problems. This includes phase transitions and some understanding of solution geometry. This lecture is based in part on joint works with Zsolt Bartha, Jian Ding, Allan Sly, and Yumeng Zhang; I will also survey related results in the area.

Nike Sun
Massachusetts Institute of Technology, U.S.
nsun@mit.edu

CP1

Engineering Kernelization for Maximum Cut

Kernelization is a general theoretical framework for preprocessing instances of NP-hard problems into (generally smaller) instances with bounded size, via the repeated application of data reduction rules. For the fundamental Max Cut problem, kernelization algorithms are theoretically highly efficient for various parameterizations. However, the efficacy of these reduction rules in practice—to aid solving highly challenging benchmark instances to optimality—remains entirely unexplored. We engineer a new suite of efficient data reduction rules that subsume most of the previously published rules, and demonstrate their significant impact on benchmark data sets, including synthetic instances, and data sets from the VLSI and image segmentation application domains. Our experiments reveal that current state-of-the-art solvers can be sped up by up to *multiple orders of magnitude* when combined with our data reduction rules. On social and biological networks in particular, kernelization enables us to solve four instances that were previously unsolved in a ten-hour time limit with state-of-the-art solvers; three of these instances are now solved in less than two seconds.

Damir Ferizovic, Demian Hespe, Sebastian Lamm
Karlsruhe Institute of Technology
dferizovic@ira.uni-karlsruhe.de, hespe@kit.edu, lamm@kit.edu

Matthias Mnich
Maastricht University / Universit{ä}t Bonn
Department of Quantitative Economics / Computer Science
mmnich@uni-bonn.de

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

Darren Strash
Department of Computer Science
Colgate University
dstrash@hamilton.edu

CP1

Shared-Memory Branch-and-Reduce for Multiterminal Cuts

We introduce the fastest known exact algorithm for the multiterminal cut problem with k terminals. In particular, we engineer existing as well as new data reduction rules. We use the rules within a branch-and-reduce framework and to boost the performance of an ILP formulation. Our algorithms achieve improvements in running time of up to multiple orders of magnitudes over the ILP formulation without data reductions, which has been the de facto standard used by practitioners. This allows us to solve instances to optimality that are significantly larger than was previously possible.

Monika Henzinger
University of Vienna, Austria
monika.henzinger@univie.ac.at

Alexander Noe
University of Vienna
alexander.noe@univie.ac.at

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

CP1

Computing Optimal Hypertree Decompositions

We propose a new algorithmic method for computing the hypertree width of hypergraphs, and we evaluate its performance empirically. At the core of our approach lies a novel ordering based characterization of hypertree width which lends to an efficient encoding to SAT modulo Theory (SMT). We tested our algorithm on an extensive benchmark set consisting of real-world instances from various sources. Our approach outperforms state-of-the-art algorithms for hypertree width. We achieve a further speedup by a new technique that first solves a relaxation of the problem and subsequently uses the solution to guide the algorithm for solving the problem itself.

Andre Schidler, Stefan Szeider
TU Wien
andre@schidler.at, stefan@szeider.net

CP1

FixCon: A Generic Solver for Fixed-Cardinality Subgraph Problems

In fixed-cardinality optimization problems in graphs, we are given a graph $G = (V, E)$, an objective function f, and an integer k and search for a set $S \subseteq V$ of k vertices that maximizes $f(G[S])$ where $G[S]$ is the subgraph of G induced by S. We implement an enumeration-based algorithm for solving fixed-cardinality optimization problems when $G[S]$ needs to be connected. To avoid enumerating all connected subgraphs of order k, we present several generic pruning rules and a generic heuristic for computing a lower bound for the objective value. We perform an experimental analysis of the performance of the algorithm and the usefulness of the pruning rules for eight example problems in which one aims to find dense, sparse, or degree-constrained connected subgraphs, respectively. Our experiments show that, when this generic solver is combined with problem-specific pruning rules, our algorithm is competitive with out-of-the-box ILP formulations for these problems.

Christian Komusiewicz, Frank Sommer
Philipps-Universität Marburg
komusiewicz@mathematik.uni-marburg.de,
fsommer@Mathematik.Uni-Marburg.de

CP2

An Almost 2-Approximation for All-Pairs of Shortest Paths in Subquadratic Time

Let $G = (V, E)$ be an unweighted undirected graph with n vertices and m edges. Dor, Halperin, and Zwick [FOCS 1996, SICOMP 2000] presented an $\tilde{O}(n^2)$-time algorithm that computes estimated distances with a multiplicative approximation of 3. Berman and Kasiviswanathan [WADS 2007] improved the approximation of Dor et al. and presented an $\tilde{O}(n^2)$-time algorithm that produces for every $u, v \in V$ an estimate $\hat{d}(u, v)$ such that: $d_G(u, v) \leq \hat{d}(u, v) \leq 2d_G(u, v) + 1$. In this paper we present an algorithm that for every $\epsilon \in (0, 1/2)$ computes in $\tilde{O}(m) + n^{2-\Omega(\epsilon)}$ time an $\tilde{O}(n^{1\frac{5}{6}})$-space data structure that in $O(1/\epsilon)$ time reports, for every $u, v \in V$, an estimate $\hat{d}(u, v)$ such that:

$$d_G(u, v) \leq \hat{d}(u, v) \leq 2(1 + \epsilon)d_G(u, v) + 5.$$

Our result improves, simultaneously, the running time and the multiplicative approximation of the $\tilde{O}(n^2)$-time $(3, 0)$-approximation algorithm of Dor et al. at the cost of introducing also an additive approximation.

Liam Roditty
Bar-Ilan University
liamr@macs.biu.ac.il

Maor Akav
BIU
maorakav@gmail.com

CP2

Tight Running Time Lower Bounds for Strong Inapproximability of Maximum K-Coverage, Unique Set Cover and Related Problems (via T-Wise Agreement Testing Theorem)

We show, assuming the randomized Gap Exponential Time Hypothesis, that the following cannot be done in $T(k)N^{o(k)}$-time for any T where N is the input size: $(1 - 1/e + \nu)$-approximation for Max k-Coverage, $(1 + 2/e - \nu)$-approximation for k-Median and $(1 + 8/e - \nu)$-approximation for k-Mean for any $\nu > 0$, any constant factor approximation for k-Unique Set Cover, k-Nearest Codeword Problem and k-Closest Vector Problem, $(1 + \mu)$-approximation for k-Minimum Distance Problem and k-Shortest Vector Problem for some $\mu > 0$. Since these problems are solvable in $N^{O(k)}$ time, our time lower bounds are essentially tight. Moreover, Max k-Coverage, k-Median and k-Mean admit polynomial or FPT time approximation with ratios $(1-1/e)$, $(1+2/e)$ and $(1+8/e)$; hence, our inapproximability ratios are also tight here. At the heart of our results, we show that Label Cover cannot be approximated to within any constant factor in $T(k)N^{o(k)}$ time where k is the number of nodes on the larger-alphabet side. This hardness of Label Cover is shown via a t-wise agreement testing theorem of the following form: given local functions $f_1, \ldots, f_k$ on domains $S_1, \ldots, S_k \subseteq [n]$, if random t functions "weakly agree' with some probability, then there is a global function g on domain $[n]$ that "mostly agrees' with many f_i's. We prove such a statement for "random-looking' sets $S_1, \ldots, S_k$ of size $O(n/k)$.

Pasin Manurangsi
UC Berkeley
pasin@berkeley.edu

CP2

Equivalences Between Triangle and Range Query Problems

We define a natural class of range query problems, and prove that all problems within this class have the same time complexity (up to polylogarithmic factors). The equivalence is very general, and even applies to online algorithms. This allows us to obtain new improved algorithms for all of the problems in the class. We then focus on the special case of the problems when the queries are offline and the number of queries is linear. We show that our range query problems are runtime-equivalent (up to polylogarithmic factors) to counting for each edge e in an m-edge graph the number of triangles through e. This natural triangle problem can be solved using the best known triangle counting algorithm, running in $O(m^{2\omega/(\omega+1)}) \leq O(m^{1.41})$ time. Moreover, if $\omega = 2$, the $O(m^{2\omega/(\omega+1)})$ running time is known to be tight (within $m^{o(1)}$ factors) under the 3SUM Hypothesis. In this case, our equivalence settles the complexity of the range query problems. Our problems constitute the first equivalence class with this peculiar running time bound. To better understand the complexity of these problems, we also provide a deeper insight into the family of triangle problems, in particular showing black-box reductions between triangle listing and per-edge triangle detection and counting. Finally, we give some not necessarily tight, but still surprising reductions from variants of matrix products, such as the (min, max)-product.

Lech Duraj, Krzysztof Kleiner, Adam Polak
Jagiellonian University
duraj@tcs.uj.edu.pl, krzysztof.kleiner@gmail.com, polak@tcs.uj.edu.pl

Virginia Vassilevska Williams
MIT
virgi@mit.edu

CP2

New Algorithms and Lower Bounds for All-Pairs Max-Flow in Undirected Graphs

We investigate the time-complexity of the All-Pairs Max-Flow problem: Given a graph with n nodes and m edges, compute for all pairs of nodes the maximum-flow value between them. If (single-pair) Max-Flow can be solved in time $T(m)$, then an $O(n^2) \cdot T(m)$ is a trivial upper bound. But can we do better? For directed graphs, recent results suggest that this time bound is essentially optimal. In contrast, for undirected graphs with edge capacities, Gomory and Hu (1961) showed a much faster algorithm, in time $O(n) \cdot T(m)$. Under the plausible assumption that Max-Flow can be solved in time $m^{1+o(1)}$, this half-century old algorithm yields an $nm^{1+o(1)}$ bound. Several other algorithms have been designed through the years, including $\tilde{O}(mn)$ time for unit-capacity edges (unconditionally), but none of them break the $O(mn)$ barrier. Meanwhile, no super-linear lower bound was shown for undirected graphs. We design the first hardness reductions for All-Pairs Max-Flow in undirected graphs, giving an essentially optimal lower bound for the node-capacities setting. For unit edge capacities, we show a new algorithm that breaks the $O(mn)$! Assuming $T(m) = m^{1+o(1)}$, our algorithm runs in time $m^{3/2+o(1)}$. Even with current Max-Flow algorithms we improve state-of-the-art as long as $m = O(n^{5/3-\varepsilon})$. Finally, we show a new $\tilde{O}(m)$ time nondeterministic algorithm for constructing a cut-equivalent tree, thus proving a non-reducibility result.

Ohad Trabelsi
Weizmann Institute of Science
ohad.trabelsi@weizmann.ac.il

Amir Abboud
IBM Almaden Research Center
amir.abboud@ibm.com

Robert Krauthgamer
Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il

CP2

Truly Subcubic Min-Plus Product for Less Structured Matrices, with Applications

The goal of this paper is to get truly subcubic algorithms for Min-Plus product for less structured inputs than what was previously known, and to apply them to versions of All-Pairs Shortest Paths (APSP) and other problems. The results are as follows: (1) Our main result is the first truly subcubic algorithm for the Min-Plus product of two $n \times n$ matrices A and B with polylog(n) bit integer entries, where B has a partitioning into $n^\epsilon \times n^\epsilon$ blocks (for any $\epsilon > 0$) where each block is at most n^δ-far (for $\delta < 3 - \omega$, where $2 \leq \omega < 2.373$) in ℓ_∞ norm from a constant rank integer matrix. (2) The first application of our main result is a truly subcubic algorithm for APSP in a new type of geometric graph, which extends the result of Chan'10 in the case of integer edge weights. (3) In the second application we consider a batch range mode problem in which one is given a length n sequence and n contiguous subsequences, and one is asked to compute the range mode of each subsequence. We give the first $O(n^{1.5-\epsilon})$ time for $\epsilon > 0$ algorithm for this problem. (4) Our final application is to the Maximum Subarray problem: given an $n \times n$ integer matrix, find the contiguous subarray of maximum entry sum. We show that Maximum Subarray can be solved in $O(n^{3-\epsilon})$ (for $\epsilon > 0$) time, as long as the entries are no larger than $O(n^{0.62})$ in absolute value.

Yinzhan Xu, Virginia Williams
Massachusetts Institute of Technology
xyzhan@mit.edu, virgi@mit.edu

CP3

Learning from Satisfying Assignments under Continuous Distributions

What kinds of functions are learnable from their satisfying assignments? Motivated by this simple question, we extend the framework of De, Diakonikolas, and Servedio [DDS15], which studied the learnability of probability distributions over $\{0,1\}^n$ defined by the set of satisfying assignments to "low-complexity" Boolean functions, to Boolean-valued functions defined over continuous domains. In our learning scenario there is a known "background distribution" D over $\mathbb{R}^n$ (such as a known normal distribution or a known log-concave distribution) and the learner is given i.i.d. samples drawn from a target distribution D_f, where D_f is D restricted to the satisfying assignments of an unknown low-complexity Boolean-valued function f. The problem is to learn an approximation D' of the target distribution D_f which has small error as measured in total variation distance. We give a range of efficient algorithms and hardness results for this problem, focusing on the case when f is a low-degree polynomial threshold function (PTF). When the background distribution D is log-concave, we show that this learning problem is efficiently solvable for degree-1 PTFs (i.e., linear threshold functions) but not for degree-2 PTFs. In contrast, when D is a normal distribution, we show that this learning problem is efficiently solvable for degree-2 PTFs but not for degree-4 PTFs. Our hardness results rely on standard assumptions about secure signature schemes.

Clément Canonne
Columbia University
ccanonne@cs.columbia.edu

Anindya De
University of Pennsylvania
anindyad@cis.upenn.edu

Rocco A. Servedio
Columbia University
rocco@cs.columbia.edu

CP3

Finding a Latent $k-$simplex in $O^*(k \cdot \texttt{Nnz}(\texttt{data}))$ Time Via Subset Smoothing

In this paper we show that learning a large class of Latent variable models, such as Mixed Membership Stochastic Block Models, Topic Models, and Adversarial Clustering etc, can be posed as special instances of the following geometric problem: *find a latent $k-$ vertex simplex, $K \subset \mathbf{R}^d$, given n data points, each obtained by perturbing a latent point in K.* This problem is not addressed in the literature and the most important contribution of this paper is to show that it admits an efficient algorithm under certain deterministic assumptions which naturally hold for the models considered here. The key insight behind the algorithm is the observation that K is close to the *subset smoothed polytope*, K', the convex hull of the $\binom{n}{\delta n}$ points obtained by averaging all δn subsets of the data points, for a given $\delta \in (0,1)$. The algorithm exploits an $O(\text{nnz})$ time optimization oracle for K' and achieves a runtime of $O^*(k\ \text{nnz})$, an input-sparsity based running time bound. The consequences of our algorithm are:

- MMSB Models and Topic Models: the first quasi-input-sparsity time algorithm for parameter estimation for $k \in O^*(1)$
- Adversarial Clustering: In $k-$means, if, an adversary is allowed to move many data points from each cluster our algorithm still estimates cluster centers well.

Chiranjib Bhattacharyya
Indian Institute of Science, Bangalore
India - 560012
chiru@iisc.ac.in

Ravindran Kannan
Indian. Institute of Science
Microsoft Research India
kannan@microsoft.com

CP3

A Tight Analysis of Greedy Yields Subexponential Time Approximation for Uniform Decision Tree

Decision Tree (DT) is a classic formulation of active learning: given n hypotheses with nonnegative weights summing to 1 and a set of tests that each partition the hypotheses, output a decision tree using the given tests that uniquely identifies each hypothesis and has minimum (weighted) average depth. Previous works showed that the greedy algorithm achieves a $O(\log n)$ approximation for this problem and it is NP-hard beat a $O(\log n)$ approximation, settling the complexity. However, for Uniform Decision Tree (UDT), i.e. DT with uniform weights, the story is more subtle. The greedy algorithm's $O(\log n)$ approximation was the best known, but the largest approximation ratio known to be NP-hard is $4-\varepsilon$. We prove that the greedy algorithm gives a $O(\frac{\log n}{\log OPT})$ approximation for UDT, where OPT is the cost of the optimal tree, and show this is best possible for the greedy algorithm. As a corollary, we resolve a conjecture of [KPB99]. Using this result, we give for all $\alpha \in (0,1)$ a $\frac{9.01}{\alpha}$ approximation algorithm of UDT using subexponential time $2^{\tilde{O}(n^\alpha)}$. As a corollary, getting a super-constant approximation ratio on UDT is not NP-hard, assuming the Exponential Time Hypothesis. This work thus adds approximating UDT to a small list of natural problems with subexponential time algorithms but no known polynomial time algorithms. We show similar results for instances of DT whose weights are not far from uniform.

Ray Li, Percy Liang
Stanford University
rayyli@stanford.edu, pliang@cs.stanford.edu

Steve Mussmann

Stanford University, U.S.
mussmann@cs.stanford.edu

CP3

List Decodable Learning Via Sum of Squares

In the list-decodable linear regression problem, we are given labelled examples $\{(X_i, y_i)\}_{i\in[N]}$ containing a subset S of βN *inliers* $\{X_i\}_{i\in S}$ that are drawn i.i.d. from standard Gaussian distribution $N(0, I)$ in $\mathbb{R}^d$, where the corresponding labels y_i are well-approximated by a linear function $\hat{\ell}$. The remaining $(1-\beta)N$ points are 'outliers' and can be adversarially chosen. We devise an algorithm that outputs a list $\mathcal{L}$ of linear functions such that there exists some $\ell \in \mathcal{L}$ that is close to $\hat{\ell}$. This yields the first efficient algorithm for linear regression in a list-decodable setting. Our results hold for any distribution of examples whose concentration and anti-concentration can be certified by Sum-of-Squares proofs.

Morris M. Yau
University of California Berkeley
1995
morrisyau@berkeley.edu

Prasad Raghavendra
U C Berkeley
prasad@cs.berkeley.edu

CP3

Oblivious Sketching of High-Degree Polynomial Kernels

Kernel methods are fundamental tools in machine learning that allow detection of non-linear dependencies between data without explicitly constructing feature vectors in high dimensional spaces. A major disadvantage of kernel methods is their poor scalability: primitives such as kernel PCA or kernel ridge regression generally take large quadratic space and (at least) quadratic time. Some methods for speeding up kernel linear algebra are known, but they all invariably take time exponential in either the dimension of the input point set (e.g., fast multipole methods suffer from *curse of dimensionality*) or in the degree of the kernel function. *Oblivious sketching* has emerged as a powerful approach to speeding up numerical linear algebra over the past decade, but our understanding of oblivious sketching for kernel matrices has remained quite limited, suffering from the aforementioned exponential dependence on input parameters. Our main contribution is a general method for applying sketching solutions developed in numerical linear algebra over the past decade to a tensoring of data points without forming the tensoring explicitly. This leads to the first oblivious sketch for the polynomial kernel with a target dimension that is only polynomially dependent on the degree of the kernel function, as well as the first oblivious sketch for the Gaussian kernel on bounded datasets that does not suffer from an exponential dependence on the dimensionality of input data.

Thomas Dybdahl Ahle
ITU
thdy@itu.dk

Michael Kapralov
EPFL
michael.kapralov@epfl.ch

Jakob B. T. Knudsen
University of Copenhagen
jakn@di.ku.dk

Rasmus Pagh
IT University of Copenhagen
pagh@itu.dk

Ameya Velingker
Google AI
ameya.velingker@epfl.ch

David Woodruff
Carnegie Mellon University
dwoodruf@cs.cmu.edu

Amir Zandieh
EPFL
amir.zandieh@epfl.ch

CP4

Optimal Space-Depth Trade-Off of Cnot Circuits in Quantum Logic Synthesis

Due to the decoherence of the state-of-the-art physical implementations of quantum computers, it is essential to parallelize the quantum circuits to reduce their depth. Two decades ago, Moore and Nilsson demonstrated that additional qubits (or *ancillae*) could be used to design "shallow' parallel circuits for quantum operators. They proved that any n-qubit CNOT circuit could be parallelized to $O(\log n)$ depth, with $O(n^2)$ ancillae. However, the near-term quantum technologies can only support limited amount of qubits, making *space-depth trade-off* a fundamental research subject for quantum-circuit synthesis. In this work, we establish an asymptotically optimal space-depth trade-off for the design of CNOT circuits. We prove that for any $m \geq 0$, any n-qubit CNOT circuit can be parallelized to $O(\max\{\log n, \frac{n^2}{(n+m)\log(n+m)}\})$ depth, with m ancillae. We show that this bound is tight by a counting argument, and further show that even with arbitrary two-qubit quantum gates to approximate CNOT circuits, the depth lower bound still meets our construction, illustrating the robustness of our result. Our work improves upon two previous results, one by Moore and Nilsson for $O(\log n)$-depth quantum synthesis, and one by Patel, Markov, and Hayes for $m = 0$. Our results can be directly extended to stabilizer circuits using an earlier result by Aaronson and Gottesman.

Jiaqing Jiang, Xiaoming Sun
Institute of Computing Technology
Chinese Academy of Sciences
jiangjiaqing@ict.ac.cn, sunxiaoming@ict.ac.cn

Shang-Hua Teng
USC
shanghua@usc.edu

Bujiao Wu
Institute of Computing Technology
Chinese Academy of Science
wubujiao@ict.ac.cn

Kewen Wu
School of Electronics Engineering and Computer Science
Peking University
kevin@pku.edu.cn

Jialin Zhang
Institute of Computing Technology
Chinese Academy of Sciences
zhangjialin@ict.ac.cn

CP4

The Directed Flat Wall Theorem

At the core of the Robertson-Seymour theory of Graph Minors lies a powerful structure theorem which captures, for any fixed graph H, the common structural features of all the graphs not containing H as a minor [N. Robertson and P. Seymour. Graph minors XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, **89** 77:1–27, 1999]. An important step towards this structure theorem is the Flat Wall Theorem [N. Robertson and P. Seymour. Graph minors XIII. The disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63:65–110, 1995], which has a lot of algorithmic applications (for example, the minor-testing and the disjoint paths problem with fixed number terminals). In this paper, we prove the *directed* analogue of this Flat Wall Theorem. Our result builds on the recent Directed Grid Theorem by two of the authors (Kawarabayashi and Kreutzer), and we hope that this is an important and significant step toward the directed structure theorem, as with the case for the undirected graph for the graph minor project.

Archontia Giannopoulou
Department of Informatics and Telecommunications
National and Kapodistrian University of Athens
archontia.giannopoulou@gmail.com

Ken-ichi Kawarabayashi
National Institute of Informatics, Japan
k_keniti@nii.ac.jp

Stephan Kreutzer
Technical University Berlin
stephan.kreutzer@tu-berlin.de

O-Joung Kwon
Incheon National University
ojoungkwon@gmail.com

CP4

Zeros of Ferromagnetic Two-Spin Systems

We study zeros of the partition functions of ferromagnetic 2-state spin systems in terms of the external field, and obtain new zero-free regions of these systems via a refinement of Asano's and Ruelle's contraction method. The strength of our results is that they do not depend on the maximum degree of the underlying graph. Via Barvinok's method, we also obtain new efficient and deterministic approximate counting algorithms. In certain regimes, our algorithm outperforms all other methods such as Markov chain Monte Carlo and correlation decay.

Heng Guo
University of Edinburgh
hguo@inf.ed.ac.uk

Jingcheng Liu
Caltech
liuexp@berkeley.edu

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

CP4

Adaptive Quantum Simulated Annealing for Bayesian Inference and Estimating Partition Functions

Markov chain Monte Carlo algorithms have important applications in counting problems and machine learning, settings that involve estimating quantities that are difficult to compute exactly. How much can quantum computers speed up classical Markov chain algorithms? In this work we consider the problem of speeding up simulated annealing algorithms, where the stationary distributions of the Markov chains are Gibbs distributions at temperatures specified according to an annealing schedule. We introduce a quantum algorithm that both adaptively constructs an annealing schedule and quantum samples at each temperature. Our adaptive schedule roughly matches the length of the best classical adaptive schedules and improves on nonadaptive schedules by roughly a quadratic factor. Our dependence on the Markov chain gap matches other quantum algorithms and is quadratically better than what classical Markov chains achieve. Our algorithm is the first to combine both of these quadratic improvements. Like other quantum walk algorithms, it also improves on classical algorithms by producing "qsamples" instead of classical samples. In constructing the annealing schedule we make use of amplitude estimation, and we show how to make amplitude estimation nondestructive at almost no additional cost, a result that may have independent interest. Finally we demonstrate how this quantum simulated annealing algorithm can be applied to partition function estimation and Bayesian inference.

Aram Harrow, Annie Y. Wei
MIT
aram@mit.edu, anniewei@mit.edu

CP4

Normalizers and Permutational Isomorphisms in Simply-Exponential Time

We show that normalizers and permutational isomorphisms of permutation groups given by generating sets can be computed in time simply exponential in the degree of the groups. The result is obtained by exploiting canonical forms for permutation groups (up to permutational isomorphism).

Daniel Wiebking
RWTH Aachen University
wiebking@informatik.rwth-aachen.de

CP5

Puzzling Grid Embeddings

We present a pipeline that, given a weighted graph as an input, produces a planar grid embedding where all edges are represented as axis-aligned straight lines with their Euclidean length matching their edge weight (if such an embedding exists). Being able to compute such embeddings is important for visualization purposes but is additionally helpful to solve certain optimization problems faster, as e.g. the Steiner tree problem. Our embedding pipeline consists of three main steps: In the first step, we identify rigid substructures which we call puzzle pieces. In the second step, we merge puzzle pieces if possible. In the third and last step, we compute the final embedding (or decide that such an embedding does not exist) via backtracking. We describe suitable data structures and engineering techniques for accelerating all steps of the pipeline along the way. Experiments on a large variety of input graphs demonstrate the applicability and scalability of our approach.

Moritz Beck, Sabine Storandt
Universität Konstanz
moritz.beck@uni-konstanz.de, sabine.storandt@uni-konstanz.de

CP5

Approximating Vertex Cover Using Structural Rounding

In this work, we provide the first practical evaluation of the structural rounding framework for approximation algorithms. Structural rounding works by first editing to a well-structured class, efficiently solving the edited instance, and "lifting" the partial solution to recover an approximation on the input. We focus on the well-studied Vertex Cover problem, and edit to the class of bipartite graphs (where Vertex Cover has an exact polynomial time algorithm). In addition to the naive lifting strategy for Vertex Cover described by Demaine et al., we introduce a suite of new lifting strategies and measure their effectiveness on a large corpus of synthetic graphs. We find that in this setting, structural rounding significantly outperforms standard 2-approximations. Further, simpler lifting strategies are extremely competitive with the more sophisticated approaches. The implementations are available as an open-source Python package, and all experiments are replicable.

Andrew van Der Poel
NC State University
ajvande4@ncsu.edu

Brian Lavallee
University of Utah
brian.lavallee@utah.edu

Hayley Russell
NC State University
hayleysrussell@gmail.com

Blair Sullivan
University of Utah
sullivan@cs.utah.edu

CP5

A Multi-Criteria Approximation Algorithm for Influence Maximization with Probabilistic Guarantees

The well-studied influence maximization problem involves choosing a seed set of a given size, which maximizes the expected influence. However, such solutions might have a significant probability of achieving low influence, which might not be suitable in many applications. In this paper, we consider a different approach: find a seed set that maximizes the influence set size with a given probability. We show that this objective is not submodular, and design a greedy, multi-criteria approximation algorithm for this problem with rigorous approximation guarantees. We also evaluate our algorithm on multiple datasets, and show that they have similar or better quality as the ones optimizing the expected influence, but with additional guarantees on the probability.

Nguyen D. Pham
University of Houston
aphamdn@gmail.com

Gopal Pandurangan
Department of Computer Science
University of Houston
gopalpandurangan@gmail.com

Anil Vullikanti
University of Virginia
vsakumar@virginia.edu

Maleq Khan
Texas A&M University Kingsville
maleq.khan@tamuk.edu

Qin Zhang
Indiana University Bloomington
qzhangcs@indiana.edu

CP5

Group Centrality Maximization for Large-Scale Graphs

The study of vertex centrality measures is a key aspect of network analysis. Naturally, such measures have been generalized to groups of vertices; for popular measures it was shown that the problem of finding the most central group is NP-hard. As a result, approximation algorithms to maximize group centralities were introduced recently. Despite a nearly-linear running time, approximation algorithms for popular group centralities are rather slow due to high constant overheads. That is why we introduce GED-Walk, a new submodular group centrality measure inspired by Katz centrality. In contrast to closeness and betweenness, it considers walks of any length rather than shortest paths, with shorter walks having a higher contribution. We define algorithms that (i) efficiently approximate the GED Walk score of a given group and (ii) efficiently approximate the (proved to be NP-hard) problem of finding a group with highest GED-Walk score. Experiments on several real-world datasets show that scores obtained by GED-Walk improve performance on classification tasks that are common in graph mining. An evaluation of empirical running times demonstrates that maximizing GED-Walk (in approximation) is two orders of magnitude faster compared to group betweenness approximation and for group sizes ≤ 100 one to two orders faster than group closeness approximation. For graphs with tens of millions of edges, approximate GED-Walk maximization typically needs less than one minute.

Eugenio Angriman, Alexander Van Der Grinten
Humboldt-Universität zu Berlin
angrimae@hu-berlin.de, avdgrinten@hu-berlin.de

Aleksandar Bojchevski, Daniel Zügner, Stephan Günnemann
Technical University of Munich
bojchevs@in.tum.de, zuegnerd@in.tum.de, guennemann@in.tum.de

Henning Meyerhenke
Humboldt-Universität zu Berlin

meyerhenke@hu-berlin.de

CP6

Symmetric Polymorphisms and Efficient Decidability of Promise Csps

In the field of constraint satisfaction problems (CSP), promise CSPs are an exciting new direction of study. In a promise CSP, each constraint comes in two forms: "strict" and "weak," and in the associated decision problem one must distinguish between being able to satisfy all the strict constraints versus not being able to satisfy all the weak constraints. The most commonly cited example of a promise CSP is the approximate graph coloring problem–which has recently benefited from multiple breakthroughs [BKO19, WZ19] due to a systematic study of promise CSPs under the lens of "polymorphisms," operations that map tuples in the strict form of each constraint to a tuple in its weak form. In this work, we present a simple algorithm which in polynomial time solves the decision problem for all promise CSPs that admit infinitely many symmetric polymorphisms, that is the coordinates are permutation invariant. This generalizes previous work of the authors [BG19]. We also extend this algorithm to a more general class of block-symmetric polymorphisms. As a corollary, this single algorithm solves all polynomial-time tractable Boolean CSPs simultaneously. These results give a new perspective on Schaefer's classic theorem and shed further light on how symmetries of polymorphisms enable algorithms.

Joshua Brakensiek
Stanford University
jbrakens@stanford.edu

Venkatesan Guruswami
Carnegie Mellon University
venkatg@cs.cmu.edu

CP6

Extended Formulation Lower Bounds for Refuting Random Csps

Random constraint satisfaction problems (CSPs) such as random 3-SAT are conjectured to be computationally intractable. The average case hardness of random 3-SAT and other CSPs has broad and far-reaching implications on problems in approximation, learning theory and cryptography. In this work, we show subexponential lower bounds on the size of linear programming relaxation for refuting random instances of constraint satisfaction problems. Formally, suppose $P : \{0,1\}^k \to \{0,1\}$ is a predicate that supports a $t-1$-wise uniform distribution on its satisfying assignments. Consider the distribution of random instances of CSP P with $m = \Delta n$ constraints. We show that any linear programming extended formulation that can refute instances from this distribution with constant probability must have size at least $\Omega\left(\exp\left(\left(\frac{n^{t-2}}{\Delta^2}\right)^{\frac{1-\nu}{k}}\right)\right)$ for all $\nu > 0$. For example, this yields a lower bound of size $\exp(n^{1/3})$ for random 3-SAT with a linear number of clauses. We use the technique of *pseudocalibration* to directly obtain extended formulation lower bounds from the planted distribution. This approach bypasses the need to construct Sherali-Adams integrality gaps in proving general LP lower bounds. As a corollary, one obtains a self-contained proof of subexponential Sherali-Adams LP lowerbounds for these problems.

Jonah Brown-Cohen, Prasad Raghavendra
UC Berkeley
jonahibc@gmail.com, raghavendra@berkeley.edu

CP6

Quasi-Popular Matchings, Optimality, and Extended Formulations

Let $G = (A \cup B, E)$ be an instance of the stable marriage problem where every vertex ranks its neighbors in a strict order of preference. A matching M in G is popular if M does not lose a head-to-head election against any matching. Popular matchings are a well-studied generalization of stable matchings, introduced with the goal of enlarging the set of admissible solutions, while maintaining a certain level of fairness. Every stable matching is a min-size popular matching. Unfortunately, when there are edge costs, it is NP-hard to find a popular matching of minimum cost – even worse, the min-cost popular matching problem is hard to approximate up to any factor. Let opt be the cost of a min-cost popular matching. Our goal is to efficiently compute a matching of cost at most opt by paying the price of mildly relaxing popularity. Our main positive result is a bi-criteria algorithm that finds in polynomial time a near-popular or 'quasi-popular' matching of cost at most opt. Key to the algorithm are a number of results for certain polytopes related to matchings. In particular, we give a polynomial-size extended formulation for an integral polytope sandwiched between the popular and quasi-popular matching polytopes. We complement these results by showing that it is NP-hard to find a quasi-popular matching of minimum cost, and that both the popular and quasi-popular matching polytopes have near-exponential extension complexity.

Yuri Faenza
Columbia University
yf2414@columbia.edu

Kavitha Telikepalli
Tata Institute of Fundamental Research
kavitha@tcs.tifr.res.in

CP6

A Deterministic Linear Program Solver in Current Matrix Multiplication Time

Interior point algorithms for solving linear programs have been studied extensively for a long time [e.g. Karmarkar 1984; Lee, Sidford FOCS'14; Cohen, Lee, Song STOC'19]. For linear programs of the form $\min_{Ax=b,x\geq 0} c^\top x$ with n variables and d constraints, the generic case $d = \Omega(n)$ has recently been settled by Cohen, Lee and Song [STOC'19]. Their algorithm can solve linear programs in $\tilde{O}(n^\omega \log(n/\delta))$ expected time, where δ is the relative accuracy. This is essentially optimal as all known linear system solvers require up to $O(n^\omega)$ time for solving $Ax = b$. However, for the case of *deterministic* solvers, the best upper bound is Vaidya's 30 years old $O(n^{2.5} \log(n/\delta))$ bound [FOCS'89]. In this paper we show that one can also settle the deterministic setting by derandomizing Cohen et al.'s $\tilde{O}(n^\omega \log(n/\delta))$ time algorithm. This allows for a strict $\tilde{O}(n^\omega \log(n/\delta))$ time bound, instead of an expected one, and a simplified analysis, reducing the length of their proof of their central path method by roughly half. Derandomizing this algorithm was an open question asked in Song's

PhD Thesis.

Jan Van Den Brand
KTH Royal Institute of Technology
janvdb@kth.se

CP6

Packing Lps Are Hard to Solve Accurately, Assuming Linear Equations Are Hard

We study the complexity of approximately solving packing linear programs. In the Real RAM model, it is known how to solve packing LPs with N non-zeros in time $\tilde{O}(N/\epsilon)$. We investigate whether the ϵ dependence in the running time can be improved. Our first main result relates the difficulty of this problem to hardness assumptions for solving dense linear equations. We show that, in the Real RAM model, unless linear equations in matrices $n \times n$ with condition number $O(n^{10})$ can be solved to ϵ accuracy faster than $\tilde{O}(n^{2.01} \log(1/\epsilon))$, no algorithm $(1-\epsilon)$-approximately solves a $O(n) \times O(n)$ packing LPs in time $\tilde{O}(n^2 \epsilon^{-0.0003})$. We currently cannot solve linear equations faster than $\tilde{O}(n^{\omega})$, where $\omega \le 2.372\ldots$ is the matrix multiplication exponent. Our second main result relates the difficulty of approximately solving packing linear programs to hardness assumptions for solving sparse linear equations: In the Real RAM model, unless Conjugate Gradient is suboptimal for solving well-conditioned sparse systems of linear equations, no algorithm $(1-\epsilon)$-approximately solves packing LPs with N non-zeros in time $\tilde{O}(N\epsilon^{-0.165})$. While we prove results in the Real RAM model, our condition number assumptions ensure that our results can be translated to fixed point arithmetic with $(\log n)^{O(1)}$ bits per number.

Rasmus Kyng
ETH Zurich
rjkyng@gmail.com

Di Wang
Google
wangd@eecs.berkeley.edu

Peng Zhang
Yale University
peng.zhang@yale.edu

CP7

Efficiency of the Floating Body As a Robust Measure of Dispersion

Among robust notions of shape, depth and dispersion of a distribution or dataset we have Tukey depth and depth curves, which are essentially the same as the convex floating body in convex geometry. These notions are important because they play the role of multidimensional quantiles and rank statistics. At the same time, they can be difficult to use because they are computationally intractable in general. We develop a theory of algorithmic efficiency for these notions for several broad and relevant families of distributions: symmetric log-concave distributions and certain multivariate stable distributions and power-law distributions. As an example of the power of these results, we show how to solve the Independent Component Analysis problem for power-law distributions, even when the first moment is infinite.

Joseph T. Anderson
Salisbury University
jtanderson@salisbury.edu

Luis Rademacher
University of California, Davis
lrademac@ucdavis.edu

CP7

Exact Computation of a Manifold Metric, Via Lipschitz Extensions and Shortest Paths on a Graph

Data-sensitive metrics adapt distances locally based the density of data points with the goal of aligning distances and some notion of similarity. In this paper, we give the first exact algorithm for computing a data-sensitive metric called the nearest neighbor metric. In fact, we prove the surprising result that a previously published 3-approximation is an exact algorithm. The nearest neighbor metric can be viewed as a special case of a density-based distance used in machine learning, or it can be seen as an example of a manifold metric. Previous computational research on such metrics despaired of computing exact distances on account of the apparent difficulty of minimizing over all continuous paths between a pair of points. We leverage the exact computation of the nearest neighbor metric to compute sparse spanners and persistent homology. We also explore the behavior of the metric built from point sets drawn from an underlying distribution and consider the more general case of inputs that are finite collections of path-connected compact sets. The main results connect several classical theories such as the conformal change of Riemannian metrics, the theory of positive definite functions of Schoenberg, and screw function theory of Schoenberg and Von Neumann. We also develop some novel proof techniques based on the combination of screw functions and Lipschitz extensions that may be of independent interest.

Gary Miller, Timothy Chu
Carnegie Mellon University
glmiller@cs.cmu.edu, tzchu@andrew.cmu.edu

Donald Sheehy
University of Connecticut
don.r.sheehy@gmail.com

CP7

Computational Concentration of Measure: Optimal Bounds, Reductions, and More

Product measures of dimension n are known to be "concentrated' under Hamming distance. More precisely, for any set S in the product space of probability $\Pr[S] \ge \epsilon$, a random point in the space, with probability $1-\delta$, has a neighbor in S that is different from the original point in only $O(\sqrt{n \cdot \ln(\frac{1}{\epsilon\delta})})$ coordinates. In this work, we obtain the tight *computational* (algorithmic) version of this result, showing how given a random point and access to an S-membership query oracle, we can find such a close point of Hamming distance $O(\sqrt{n \cdot \ln(\frac{1}{\epsilon\delta})})$ in time $poly(n, 1/\epsilon, 1/\delta)$. This resolves an open question of [Mahloujifar and Mahmoody, ALT19]. As corollaries, we obtain polynomial-time poisoning and evasion attacks against learning algorithms when the original vulnerabilities have any cryptographically non-negligible probability. Additionally, we define a new notion of algorithmic reduction between computational concentration of measure in different metric probability spaces using which we prove a

computational concentration theorem for Gaussian under ℓ_1. We further prove several extensions to the results to (1) the case of weighted Hamming distance, (2) concentration around mean, and (3) discrete random processes. We also prove an exponential lower bounds on the average running time of *non-adaptive* query algorithms.

Omid Etesami
IPM
etesami@gmail.com

Saeed Mahloujifar, Mohammad Mahmoody
University of Virginia
saeed@virginia.edu, mohammad@virginia.edu

CP7

Sample Efficient Toeplitz Covariance Estimation

How many samples from a distribution D over d-dimensional vectors are necessary to learn the distribution's covariance matrix, T? Moreover, how can we leverage a priori knowledge about T's structure to reduce this sample complexity? I will discuss this fundamental statistical problem in the setting where T is known to have Toeplitz structure. Toeplitz covariance matrices arise in countless signal processing applications, from wireless communications, to medical imaging, to time series analysis. In many of these applications, we are interested in algorithms that only view a subset of entries in each d-dimensional vector sample from D. We care about minimizing two notions of sample complexity 1) the total number of vector samples taken and 2) the number of entries accessed in each vector sample. I will present several new non-asymptotic bounds on these sample complexity measures. We will start by taking a fresh look at classical and widely used algorithms, including methods based on selecting entries from each sample according to a "sparse ruler". I then will introduce a novel sampling and estimation strategy that improves on existing methods in many settings by utilizing tools from random matrix sketching, leverage score sampling for continuous signals, and sparse Fourier transform algorithms. Our result fits into a broader line of work which seeks to address fundamental problems in signal processing using tools from theoretical computer science.

Yonina Eldar
Weizmann Institute of Science
yonina.eldar@weizmann.ac.il

Jerry Li
Microsoft Research
jerrl@microsoft.com

Cameron Musco
University of Massachusetts Amherst
cmusco@cs.umass.edu

Christopher Musco
New York University
cmusco@nyu.edu

CP7

On the Learnability of Random Deep Networks

In this paper we study the learnability of random deep networks both theoretically and experimentally. On the theoretical front, assuming the statistical query model, we show that the learnability of random deep networks with sign and sigmoid activation drops exponentially with their depths. The main point we use in the argument is that even for highly correlated inputs the outputs of deep random networks are highly un-correlated (near orthogonal) – this suggests the possibility that deep random networks may be acting as (pseudo) random functions. We verify the learnability of deep random networks experimentally and find that the learnability drops sharply with depth even with the state-of-the-art training methods.

Rina Panigrahy
Microsoft Research Silicon Valley, USA
rinapy@gmail.com

CP8

Extremal Distances in Directed Graphs: Tight Spanners and Near-Optimal Approximation Algorithms

Given a directed graph $G = (V, E)$ on n vertices and m edges, a subgraph $H = (V, E' \subseteq E)$ is defined to be a t-diameter spanner if the diameter of H is at most t times the diameter of G. We show the existence of (and algorithms to compute) various t-diameter spanners with a sparse set of edges and $t < 2$, for directed graphs. In addition, we show that our spanner constructions give tight bounds on the number of edges. To the best of our knowledge, our work is the first to focus on the existence of various sparse (with $\ll n^2$ edges) diameter spanners of stretch < 2, for directed graphs. We also study *eccentricity spanner*, which is a subgraph that approximately preserves all vertex eccentricities of the original graph. As an application of our eccentricity spanner construction, we obtain the *first* $\widetilde{O}(m)$*-time algorithm* for computing 2-approximation of vertex eccentricities in general directed graphs. This improves the result of Backurs et al. [STOC 2018] who gave an $\widetilde{O}(m\sqrt{n})$ time algorithm for this problem, and showed that there is no $O(n^{2-o(1)})$ time algorithm that achieves approximation better than 2, unless SETH fails; this shows that our approximation factor is essentially tight. Finally, we study extremal distance spanners under dynamic settings.

Keerti Choudhary
Weizmann Institute of Science
keerti.choudhary@weizmann.ac.il

Omer Gold
Tel Aviv University
omergolden@gmail.com

CP8

Shorter Labeling Schemes for Planar Graphs

An *adjacency labeling scheme* for a given class of graphs is an algorithm that for every graph G from the class, assigns bit strings (labels) to vertices of G so that for any two vertices u, v, whether u and v are adjacent can be determined by a fixed procedure that examines only their labels. It is known that planar graphs with n vertices admit a labeling scheme with labels of bit length $(2 + o(1)) \log n$. In this work we improve this bound by designing a labeling scheme with labels of bit length $(\frac{4}{3} + o(1)) \log n$. In graph-theoretical terms, this implies an explicit construction of a graph on $n^{4/3+o(1)}$ vertices that contains all planar graphs on n vertices as induced subgraphs, improving the previous best upper bound of $n^{2+o(1)}$. Our scheme generalizes to graphs of bounded Euler genus with the same label length up to a second-order term. All the labels of the input graph can be computed in polynomial time, while adjacency can be decided from the labels in constant time.

Cyril Gavoille
Bordeaux University
gavoille@labri.fr

Marthe Bonamy
CNRS-LaBRI, University of Bordeaux
marthe.bonamy@u-bordeaux.fr

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

CP8

Coarse-Grained Complexity for Dynamic Algorithms

To date, the only way to get polynomial lower bounds for dynamic algorithms is via fine-grained complexity, which relies on strong assumptions about specific problems such as the Strong Exponential Time Hypothesis and the Online Matrix-Vector Multiplication Conjecture. While they have led to many exciting discoveries, dynamic algorithms still miss out some benefits and lessons from the traditional coarse-grained approach that relates together classes of problems such as P and NP. We initiate the study of coarse-grained complexity theory for dynamic algorithms, and obtain the following results. 1. A research program for proving polynomial unconditional lower bounds for dynamic orthogonal vector (OV) in the cell-probe model is motivated by the fact that many conditional lower bounds can be shown via reductions from the dynamic OV problem. Since the cell-probe model is more powerful than word RAM and has historically allowed smaller upper bounds (e.g. [Larsen, Williams, SODA 2017]), it might turn out that dynamic OV is easy in the cell-probe model, making this research direction infeasible. Our theory implies that if this is the case, then it will have some very interesting algorithmic consequences. 2. We show that there are no efficient reductions (in both cell-probe and word RAM models) from dynamic OV to k-edge connectivity under an assumption about the classes of dynamic algorithms whose analogue in the static setting is widely believed.

Sayan Bhattacharya
University of Warwick
Coventry
S.Bhattacharya@warwick.ac.uk

Danupon Nanongkai
KTH Royal Institute of Technology, Sweden
danupon@gmail.com

Thatchaphol Saranurak
Toyota Technological Institute at Chicago
saranurak@ttic.edu

CP8

How to Store a Random Walk

Motivated by storage applications, we study the following data structure problem: An encoder wishes to store a collection of jointly-distributed files $X := (X_1, X_2, \ldots, X_n) \sim \mu$ which are correlated ($H_\mu(\overline{X}) \ll \sum_i H_\mu(X_i)$), using as little (expected) memory as possible, such that each individual file X_i can be recovered quickly with few (ideally constant) memory accesses. In the case of independent random files, a dramatic result by Dodis, Pătraşcu and Thorup (STOC'10) shows that it is possible to store $\overline{X}$ using just a constant number of extra bits beyond the information-theoretic minimum space, while at the same time decoding each X_i in constant time. However, in the (realistic) case where the files are correlated, much weaker results are known, requiring at least $\Omega(n/\lg n)$ extra bits for constant decoding time, even for "simple" joint distributions μ. We focus on the natural case of compressing Markov chains, i.e., storing a length-n random walk on any (possibly directed) graph G. Denoting by $\kappa(G, n)$ the number of length-n walks on G, we show that there is a succinct data structure storing a random walk using $\lg_2 \kappa(G, n) + O(\lg n)$ bits of space, such that any vertex along the walk can be decoded in $O(1)$ time on a word-RAM. If the graph is strongly connected (e.g., undirected), the space can be improved to only $\lg_2 \kappa(G, n) + 5$ extra bits.

Emanuele Viola
Northeastern University
viola@ccs.neu.edu

Omri Weinstein
Columbia University
omri@cs.columbia.edu

Huacheng Yu
Princeton University
yuhch123@gmail.com

CP8

Dynamic Low-Stretch Spanning Trees in Subpolynomial Time

Low-stretch spanning tree has been an important graph-theoretic object, as it is one of the building blocks for fast algorithms that solve symmetrically diagonally dominant linear systems, and a significant line of research has been devoted to finding constructions with optimal average stretch. In a very recent work by Goranci and Forster [STOC 2019], the authors initiated the study of low-stretch spanning trees in the dynamic setting, and they proposed a dynamic algorithm that maintains a spanning tree in $n^{\frac{1}{2}+o(1)}$ amortized update time with subpolynomial stretch in an unweighted graph on n vertices undergoing edge insertions and deletions demanded by an oblivious adversary. Our main results are twofold. First, we substantially improve the update time of Goranci and Forster [STOC 2019] from $n^{\frac{1}{2}+o(1)}$ to a subpolynomial of $n^{o(1)}$. Second, we generalize our result to weighted graphs under the decremental setting. As far as we know, this is the first non trivial dynamic algorithm for maintaining low-stretch spanning tree for weighted graphs.

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

Tianyi Zhang
IIIS, Tsinghua
tianyi-z16@mails.tsinghua.edu.cn

CP9

Approximating Multiobjective Shortest Path in Practice

We consider the multiobjective shortest path (MOSP)

problem. While known approximation algorithms allow a polynomial running time, the degrees of these polynomials are dependent on the number of objective functions. Unfortunately, this also holds true for their best-case. Exact algorithms, while attaining an exponential worst-case running time even in the number of nodes, allow for far better best-case performance and are thus preferred in practice. We introduce a new general approximation framework for MOSP. It aims at combining strong worst-case guarantees with practically useful performance. It allows for various labeling strategies as employed by exact algorithms; thus, decades of research can be utilized. We conduct a comprehensive computational study to compare our framework to known approximations and exact algorithms. For many, this is their first practical investigation. Furthermore, this is the first time that graphs of practically relevant sizes as well as real-world instances are considered in the context of MOSP approximation. The results show that our framework is superior to the known approximation methods in running time and quality. They also demonstrate the usefulness and limits of approximations compared to exact methods.

Fritz Bökler, Markus Chimani
Osnabrück University
fboekler@uos.de, markus.chimani@uos.de

CP9

Fully Dynamic Single-Source Reachability in Practice: An Experimental Study

Given a directed graph and a source vertex, the *fully dynamic single-source reachability* problem is to maintain the set of vertices that are reachable from the given vertex, subject to edge deletions and insertions. While there has been a large body of theoretical work on this problem, there has been no experimental study that compares the performance of fully dynamic reachability algorithms in practice. Previous experimental studies in this area concentrated only on the more general all-pairs reachability or transitive closure problem and did not use real-world dynamic graphs. In this paper, we bridge this gap by empirically studying an extensive set of algorithms for the single-source reachability problem in the fully dynamic setting. In particular, we design several fully dynamic variants of well-known approaches to obtain and maintain reachability information with respect to a distinguished source. Moreover, we extend the existing insertions-only or deletions-only upper bounds into fully dynamic algorithms. Even though the worst-case time per operation of all the fully dynamic algorithms we evaluate is at least linear in the number of edges in the graph (as is to be expected given the conditional lower bounds) we show in our extensive experimental evaluation that their performance differs greatly, both on generated as well as on real-world instances.

Kathrin Hanauer
University of Vienna
kathrin.hanauer@univie.ac.at

Monika Henzinger
University of Vienna, Austria
monika.henzinger@univie.ac.at

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

CP9

Karp-Sipser Based Kernels for Bipartite Graph Matching

We consider Karp–Sipser, a well known matching heuristic in the context of data reduction for the maximum cardinality matching problem. We describe an efficient implementation as well as modifications to reduce its time complexity in worst case instances, both in theory and in practical cases. We compare experimentally against its widely used simpler variant and show cases for which the full algorithm yields better performance.

Ioannis Panagiotas
ENS de Lyon
INRIA
ioannis.panagiotas@ens-lyon.fr

Kamer Kaya
Sabanci University,
Faculty of Engineering and Natural Sciences
kaya@sabanciuniv.edu

Johannes Langguth
Simula research laboratory
langguth@simula.no

Bora Uçar
CNRS
LIP ENS de Lyon
bora.ucar@ens-lyon.fr

CP9

Shrinking Trees Not Blossoms:A Recursive Maximum Matching Approach

We suggest new concepts for engineering maximum matching algorithms on general graphs which can be used as alternatives in augmenting path approaches such as, e.g., Edmonds or Micali and Vazirani. Our newly introduced *alternating rooted sets (ARSs)* for finding augmenting paths generalize the state-of-the-art *alternating trees.* In order to realize an ARS approach on multiple growing trees, we suggest the *cherry tree* data structure that does not only avoid explicit blossom shrinking but also supports a rotation operation. This operation allows for determining maximum sets of augmenting paths with respect to the given ARSs. These sets can be found by solving a maximum matching problem in the metagraph arising by shrinking the ARSs. We experimentally evaluate our new recursive metagraph approach on a wide set of benchmark instances including a comparison to publically available state-of-the-art software.

Andre Droschinsky, Petra Mutzel, Erik Thordsen
TU Dortmund University
andre.droschinsky@tu-dortmund.de, petra.mutzel@tu-dortmund.de, erik.thordsen@tu-dortmund.de

CP10

Individual Sensitivity Preprocessing for Data Privacy

The sensitivity metric in differential privacy, which is informally defined as the largest marginal change in output between neighboring databases, is of substantial significance in determining the accuracy of private data analyses. Techniques for improving accuracy when the average sensitivity is much smaller than the worst-case sensitivity have been developed within the differential privacy literature, including tools such as smooth sensitivity, Sample-and-Aggregate, Propose-Test-Release, and Lipschitz extensions. In this work, we provide a new and general Sensitivity-Preprocessing framework for reducing sensitivity, where efficient application gives state-of-the-art accuracy for privately outputting the important statistical metrics median and mean when no underlying assumptions are made about the database. In particular, our framework compares favorably to smooth sensitivity for privately outputting median, in terms of both running time and accuracy. Furthermore, because our framework is a preprocessing step, it can also be complementary to smooth sensitivity and any other private mechanism, where applying both can achieve further gains in accuracy. We additionally introduce a new notion of individual sensitivity and show that it is an important metric in the variant definition of personalized differential privacy. We show that our algorithm can extend to this context and serve as a useful tool for this variant definition and its applications in markets for privacy.

Rachel Cummings, David Durfee
Georgia Institute of Technology
rachelc@gatech.edu, ddurfee42@gmail.com

CP10

Differentially Private Release of Synthetic Graphs

We propose a differentially private mechanism that, given an input graph G with n vertices and m edges, in polynomial time generates an (epsilon, delta)-differentially private synthetic graph G approximating all cuts of the input graph up to an additive error of $O(\sqrt{(mn/\epsilon)}\log^2(n/\delta))$. This is the first construction of differentially private cut approximators that allows additive error o(m) for all $m > n\log^C n$. The previous best results gave additive $n^{3/2}$ error for every m and hence only retained information about the cut structure on very dense graphs. We also present lower bounds showing that our utility/privacy trade-off is essentially best possible if one seeks to get only additive cut approximations.

Marek Elias, Michael Kapralov
EPFL
marek.elias@epfl.ch, michael.kapralov@epfl.ch

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

Yin Tat Lee
Microsoft Research
yintatlee@gmail.com

CP10

Exponential Separations in Local Differential Privacy

We prove a general connection between the *communication* complexity of two-player games and the *sample* complexity of their multi-player locally private analogues. We use this connection to prove sample complexity lower bounds for locally differentially private protocols as straightforward corollaries of results from communication complexity. In particular, we 1) use a communication lower bound for the hidden layers problem to prove an exponential sample complexity separation between sequentially and fully interactive locally private protocols, and 2) use a communication lower bound for the pointer chasing problem to prove an exponential sample complexity separation between k-round and $(k+1)$-round sequentially interactive locally private protocols, for every k.

Matthew Joseph
University of Pennsylvania
Department of Computer and Information Science
majos@cis.upenn.edu

Jieming Mao
Google Research New York
maojm@google.com

Aaron Roth
University of Pennsylvania, USA
aaroth@cis.upenn.edu

CP10

The Rank of Sparse Random Matrices

We determine the rank of a random matrix over an arbitrary field with prescribed numbers of non-zero entries in each row and column. As an application we obtain a formula for the rate of low-density parity check codes. This formula vindicates a conjecture of Lelarge [Proc. IEEE Information Theory Workshop 2013]. The proofs are based on coupling arguments and a novel random perturbation, applicable to any matrix, that likely diminishes the number of short linear relations.

Amin Coja-Oghlan
Goethe University, Germany
acoghlan@math.uni-frankfurt.de

Alperen Ergür
Carnegie Mellon University
aergur@cs.cmu.edu

Pu Gao
University of Waterloo
p3gao@uwaterloo.ca

Samuel Hetterich, Maurice Rolvien
Goethe University
hetteric@math.uni-frankfurt.de, rolvien@math.uni-frankfurt.de

CP10

Locally Private K-Means Clustering

We design a new algorithm for the Euclidean k-means problem that operates in the local model of differential privacy. Unlike in the non-private literature, differentially private algorithms for the k-means incur both additive and multiplicative errors. Our algorithm significantly reduces the additive error while keeping the multiplicative error the same as in previous state-of-the-art results. Specifically, on a database of size n, our algorithm guarantees $O(1)$ multiplicative error and $\approx n^{1/2+a}$ additive error for an arbitrarily small constant a, whereas all previous algorithms in the local model had additive error $\approx n^{2/3+a}$. We give a simple lower bound showing that additive error of $\approx \sqrt{n}$ is necessary for k-means algorithms in the local model (at least for algorithms with a constant number of interaction rounds, which is the setting we consider in this paper).

Uri Stemmer
Ben Gurion University
stemmer@cs.bgu.ac.il

CP11

Faster Deterministic and Las Vegas Algorithms for Offline Approximate Nearest Neighbors in High Dimensions

We present a *deterministic*, truly subquadratic algorithm for offline $(1+\varepsilon)$-approximate nearest or farthest neighbor search (in particular, the closest pair or diameter problem) in Hamming space in any dimension $d \leq n^{\delta}$, for a sufficiently small constant $\delta > 0$. The running time of the algorithm is roughly $n^{2-\varepsilon^{1/2+O(\delta)}}$ for nearest neighbors, or $n^{2-\Omega(\sqrt{\varepsilon}/\log(1/\varepsilon))}$ for farthest. The algorithm follows from a simple combination of expander walks, Chebyshev polynomials, and rectangular matrix multiplication. We also show how to eliminate errors in the previous Monte Carlo randomized algorithm of Alman, Chan, and Williams [FOCS'16] for offline approximate nearest or farthest neighbors, and obtain a *Las Vegas* randomized algorithm with expected running time $n^{2-\Omega(\varepsilon^{1/3}/\log(1/\varepsilon))}$. Finally, we note a simplification of Alman, Chan, and Williams' method and obtain a slightly improved *Monte Carlo* randomized algorithm with running time $n^{2-\Omega(\varepsilon^{1/3}/\log^{2/3}(1/\varepsilon))}$. As one application, we obtain improved deterministic and randomized $(1+\varepsilon)$-approximation algorithms for MAX-SAT.

Josh Alman
MIT
jalman@seas.harvard.edu

Timothy M. Chan
University of Illinois at Urbana-Champaign
tmc@illinois.edu

Ryan Williams
Massachusetts Institute of Technology
rrw@mit.edu

CP11

Locally Consistent Parsing for Text Indexing in Small Space

We consider two closely related problems of text indexing in a sub-linear working space. The first problem is the Sparse Suffix Tree (SST) construction of a set of suffixes B using only $O(|B|)$ words of space. The second problem is the Longest Common Extension (LCE) problem, where for some parameter $1 \leq \tau \leq n$, the goal is to construct a data structure that uses $O(\frac{n}{\tau})$ words and can compute the longest common prefix length of any pair of suffixes. We show how to use ideas based on the Locally Consistent Parsing technique, that was introduced by Sahinalp and Vishkin [STOC'94], in some non-trivial ways in order to improve the known results for the above problems. We introduce the first Las Vegas SST construction algorithm that takes $O(n)$ time. This is an improvement over the result of Gawrychowski and Kociumaka [SODA '17] who obtained $O(n)$ time Monte Carlo algorithm, and $O(n\sqrt{\log|B|})$ time Las Vegas algorithm. In addition, we introduce a randomized Las Vegas construction for an LCE data structure that can be constructed in linear time and answers queries in $O(\tau)$ time. We introduce the first almost linear time deterministic SST construction, which takes $O(n \log \frac{n}{|B|})$ time. For the LCE problem, we introduce a data structure that answers LCE queries in $O(\tau\sqrt{\log^* n})$ time, with $O(n \log \tau)$ construction time. This data structure improves both query and construction time upon the results of Tanimura et al. [CPM'16].

Shay Golan
Bar-Ilan University
golansh1@cs.biu.ac.il

Or Birenzwige, Ely Porat
Bar Ilan University
birenzo@cs.biu.ac.il, porately@cs.biu.ac.il

CP11

A Lower Bound for Jumbled Indexing

In this paper we study lower bounds for a variant of *jumbled-indexing* problem: given an input string S on an alphabet $\Sigma = \{\sigma_1, \ldots, \sigma_\lambda\}$, store it in a data structure such that given frequencies $f_1, \cdots, f_\lambda$, one can find all the substrings S' of S where the frequency of the character σ_i is f_i, for $1 \leq i \leq \lambda$. This is a very interesting and challenging text indexing problem and it has received significant attention lately, both in the string-indexing community as well as in the theory community. It is known that when $\lambda = 2$, one can use a linear-size data structure to decide whether there exists a substring that matches the query in constant time. Amir et al. showed that for constant $\lambda \geq 3$, there exists a constant α_λ such that it is not possible to achieve both $O(n^{2-\alpha_\lambda})$ preprocessing time and $O(n^{1-\alpha_\lambda})$ query time. We study a variant of the problem where the goal is to *report* all the substrings of S that match a given jumbled-indexing query. Assuming the data structure operates in the pointer machine model, we prove *unconditional space* lower bounds that also apply to *binary* alphabets: we show that if the data structure has the query time of $O(n^{0.5-o(1)} + k)$, where k is the output size of the query, then it must consume $\Omega(n^{2-o(1)})$ space.

Peyman Afshani
Aarhus University
peyman@cs.au.dk

Ingo Van Duijn
Aalborg University
ingo@cs.aau.dk

Rasmus Killmann
Aarhus University
killmann@cs.au.dk

Jesper Sindahl Nielsen
Uber
jasn@cs.au.dk

CP11

Flushing Without Cascades

Buffer flushing is a technique for transforming standard external-memory search trees into write-optimized data structures. In exchange for faster amortized insertion, buffer flushing can sometimes significantly increase the latency of operations by causing so-called flushing cascades. In this paper, we show that flushing cascades are not a fundamental consequence of the buffer-flushing technique, and can be removed entirely using randomization techniques. The underlying implementation of buffer flushing relies on a buffer eviction strategy at each node in the tree. The ability for the user to select the buffer eviction strategy based on the workload has been shown to be important for performance, both in in theory and in practice. In order to support arbitrary buffer eviction strategies, we introduce the notion of a universal flush, which uses a universal eviction strategy that can be parameterized in order to simulate any other eviction strategy. This abstracts away the underlying eviction strategy, even allowing for workload-specific strategies that change dynamically.

Michael A. Bender, Rathish Das
Stony Brook University
bender@cs.stonybrook.edu, radas@cs.stonybrook.edu

Martin Farach-Colton
Rutgers University
martin@farach-colton.com

Rob Johnson
VMWare Research
robj@vmware.com

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

CP11

Better Data Structures for Colored Orthogonal Range Reporting

Range searching on categorical, or "colored', data has been studied extensively for over two decades. In this paper, we obtain the current best results for perhaps the most basic, and most often studied, version of the geometric problem: colored orthogonal range reporting. Given n colored points in two-dimensional space $[U]^2$, we present a data structure with $O(n \log^{3/4+\varepsilon} n)$ space, for an arbitrarily small constant $\varepsilon > 0$, so that all k distinct colors in any axis-aligned query rectangle can be reported in (optimal) $O(\log \log U + k)$ time; this is the first method to break the $O(n \log n)$ space barrier. In three dimensions, we present a data structure with $O(n \log^{9/5+\varepsilon} n)$ space and $O(\log n/ \log \log n + k)$ time; this improves the previous space bound of $O(n \log^4 n)$.

Timothy M. Chan
UIUC
tmc@illinois.edu

Yakov Nekrich
University of Waterloo
yakov.nekrich@googlemail.com

CP12

Very Fast Construction of Bounded Degree Spanning Graphs Via the Semi-Random Graph Process

We consider a recently proposed semi-random graph process, described as follows: we start with an empty graph on n vertices. In each round, a single player, called Builder, receives a uniformly random vertex v, and must immediately (in an online manner) choose another vertex u, adding the edge $\{u, v\}$ to the graph. Builder's end goal is to make the constructed graph satisfy some predetermined monotone graph property. Specifically, consider the property $\mathcal{P}_H$ of containing a spanning graph H as a subgraph. It was asked by N. Alon whether for any bounded-degree H, Builder can construct a graph satisfying $\mathcal{P}_H$ with high probability in $O(n)$ rounds. We answer this question positively in a strong sense, showing that any graph with maximum degree Δ can be constructed with high probability in $(3\Delta/2 + o(\Delta))n$ rounds, where the $o(\Delta)$ term tends to zero as $\Delta \to \infty$. This is tight (even for the offline case) up to a multiplicative factor of 3. Furthermore, for the special case where H is a spanning forest of maximum degree Δ, we show that H can be constructed with high probability in $O(n \log \Delta)$ rounds. This is tight, even for the offline setting. Finally, we show a separation between adaptive and non-adaptive strategies, proving a lower bound of $\Omega(n\sqrt{\log n})$ on the number of non-adaptive rounds necessary to eliminate all isolated vertices w.h.p., which is again tight.

Omri Ben-Eliezer, Lior Gishboliner
Tel Aviv University
omrib@mail.tau.ac.il, liorgis1@post.tau.ac.il

Dan Hefetz
Ariel University
danhe@ariel.ac.il

Michael Krivelevich
Tel Aviv University
krivelev@tauex.tau.ac.il

CP12

A Randomly Weighted Minimum Spanning Tree with a Random Cost Constraint

We study the minimum spanning tree problem on the complete graph where an edge e has a weight W_e and a cost C_e, each of which is an independent uniform $[0, 1]$ random variable. There is also a constraint that the spanning tree T must satisfy $C(T) \le c_0$. We establish the asymptotic value of the optimum weight via the consideration of a dual problem. The proof is therefore constructive i.e. can be thought of as the analysis of a polynomial time algorithm. We also study the minimum spanning arborescence problem on the complete digraph $\vec{K}_n$ where an edge e has a weight W_e and a cost C_e, each of which is an independent uniform $[0, 1]$ random variable. There is also a constraint that the spanning arborescence T must satisfy $C(T) \le c_0$. We establish the asymptotic value of the optimum weight via the consideration of a dual problem. The proof is via the analysis of a polynomial time algorithm.

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

CP12

Sandwiching Random Regular Graphs Between Binomial Random Graphs

Kim and Vu made the following conjecture (*Advances in Mathematics*, 2004): if $d \gg \log n$, then the random d-regular graph $\mathcal{G}(n, d)$ can asymptotically almost surely be "sandwiched' between $\mathcal{G}(n, p_1)$ and $\mathcal{G}(n, p_2)$ where p_1 and p_2 are both $(1 + o(1))d/n$. They proved this conjecture for $\log n \ll d \le n^{1/3-o(1)}$, with a defect in the sandwiching: $\mathcal{G}(n, d)$ contains $\mathcal{G}(n, p_1)$ perfectly, but is not completely contained in $\mathcal{G}(n, p_2)$. Recently, the embedding $\mathcal{G}(n, p_1) \subseteq \mathcal{G}(n, d)$ was improved by Dudek, Frieze, Ruciński and Šileikis to $d = o(n)$. In this paper, we prove Kim–Vu's sandwich conjecture, with perfect containment on both sides, for all $d \gg n/\sqrt{\log n}$. For $d = O(n/\sqrt{\log n})$, we prove a weaker version of the sandwich conjecture with p_2 approximately equal to $(d/n)\log n$, without any defect. In addition to sandwiching regular graphs, our results cover graphs whose degrees are asymptotically equal. As applications, we obtain new results on the properties of random graphs with given near-regular degree sequences, including Hamiltonicity and universality in subgraph containment, chromatic number, small subgraph counts, diameter, and many phase transitions in edge percolation.

Brendan McKay
Australian National University
brendan.mckay@anu.edu.au

Mikhail Isaev
Monash University
mikhail.isaev@monash.edu

Pu Gao
University of Waterloo
pu.gao@uwaterloo.ca

CP12

Factors and Loose Hamilton Cycles in Sparse Pseudo-Random Hypergraphs

We investigate the emergence of spanning structures in sparse pseudo-random k-uniform hypergraphs, using the following comparatively weak notion of pseudo-randomness. A k-uniform hypergraph H on n vertices is called (p, α, ϵ)-*pseudo-random* if for all not necessarily disjoint sets $A_1, \ldots, A_k \subset V(H)$ with $|A_1| \cdots |A_k| > \alpha n^k$ we have $e(A_1, \ldots, A_k) = (1 \pm \epsilon)p|A_1| \cdots |A_k|$. For any linear k-uniform F we provide a bound on $\alpha = \alpha(n)$ in terms of $p = p(n)$ and F, such that (under natural divisibility assumptions on n) any $(p, \alpha, o(1))$-pseudo-random n-vertex H with a mild minimum degree condition contains an F-factor. The approach also enables us to investigate the existence of loose Hamilton cycles. All results imply corresponding bounds for other notions of hypergraph pseudo-randomness such as jumbledness and small second eigenvalue. As a consequence of our results, perfect matchings appear at $\alpha = o(p^k)$ while loose Hamilton cycles appear at $\alpha = o(p^{k-1})$. This extends the works of Lenz–Mubayi, and Lenz–Mubayi–Mycroft who studied the analogous problems in the dense setting.

Jie Han
Department of Mathematics
University of Rhode Island
jie_han@uri.edu

Hiep Han
Universidad de Santiago de Chile
hiep.han@usach.cl

Patrick Morris
Freie Universitat Berlin
pm0041@mi.fu-berlin.de

CP12

A New Algorithm for the Robust Semi-Random Independent Set Problem

We study the independent set problem in a semi-random model proposed by Feige and Kilian. This model selects a graph with a planted independent set of size k and then allows an adversary to modify a large fraction of edges: the subgraph induced by the complement of the independent set can be modified arbitrarily, and the adversary may add (but not delete) edges from the independent set to its complement. In particular, the adversary can create a graph in which the initial planted independent set is not the largest independent set. Feige and Kilian presented a randomized algorithm, which with high probability recovers an independent set of size at least k (which may not be the planted one) when $k = \alpha n$ where α is a constant, and the probability of a random edge $p > (1 + \epsilon)\ln n/\alpha n$. We give a new deterministic algorithm in the Feige-Kilian model that finds an independent set of size at least $.99k$ provided that the planted set has size $k = \Omega(n^{2/3}/p^{1/3})$, and finds a list of independent sets, one of which is the planted one provided that $k = \Omega(n^{2/3}/p)$. This improves on the algorithm of Feige and Kilian by working for smaller k if $p = \Omega(1/n^{1/3})$, and improves on an algorithm of Steinhardt by working for slightly smaller k and by working against a stronger adversarial model. The ability to find a good approximation of the largest independent set is new when $p < \ln n/k$.

Theo R. McKenzie, Hermish Mehta
University of California, Berkeley
mckenzie@math.berkeley.edu, hermish@berkeley.edu

Luca Trevisan
Bocconi University
l.trevisan@unibocconi.it

CP13

Engineering Top-Down Weight-Balanced Trees

Weight-balanced trees are a popular form of self-balancing binary search trees. Their popularity is due to desirable guarantees, for example regarding the required work to balance annotated trees. While usual weight-balanced trees perform their balancing operations in a bottom-up fashion after a modification to the tree is completed, there exists a top-down variant which performs these balancing operations during descend. This variant has so far received only little attention. We provide an in-depth analysis and engineering of these top-down weight-balanced trees, demonstrating their superior performance. We also gaining insights into how the balancing parameters necessary for a weight-balanced tree should be chosen — with the surprising observation that it is often beneficial to choose parameters which are not feasible in the sense of the correctness proofs for the rebalancing algorithm.

Lukas Barth
Karlsruhe Institute of Technology
lukas.barth@kit.edu

Dorothea Wagner
Karlsruhe Institute of Technology
Institute of Theoretical Informatics
dorothea.wagner@kit.edu

CP13

Reverse-Safe Data Structures for Text Indexing

We introduce the notion of reverse-safe data structures.

These are data structures that prevent the reconstruction of the data they encode (i.e. they cannot be easily reversed). A data structure D is called *z-reverse-safe* when there exist at least z datasets with the same set of answers as the ones stored by D. The main challenge is to ensure that D stores as many answers to useful queries as possible, is constructed efficiently, and has size close to the size of the original dataset it encodes. Given a text of length n and an integer z, we propose an algorithm which constructs a z-reverse-safe data structure that has size $O(n)$ and answers pattern matching queries of length at most d optimally, where d is maximal for any such z-reverse-safe data structure. The construction algorithm takes $O(n^\omega \log d)$ time, where ω is the matrix multiplication exponent. We show that, despite the n^ω factor, our engineered implementation takes only a few minutes to finish for million-letter texts. We further show that plugging our method in data analysis applications gives insignificant or no data utility loss. Finally, we show how our technique can be extended to support applications under a realistic adversary model.

Giulia Bernardini
University of Milan - Bicocca
giulia.bernardini@unimib.it

Huiping Chen
King's College London
huiping.chen@kcl.ac.uk

Gabriele Fici
Università degli Studi di Palermo
gabriele.fici@unipa.it

Grigorios Loukides
King's College London
grigorios,loukides@kcl.ac.uk

Solon P. Pissis
CWI Amsterdam
solon.pissis@cwi.nl

CP13

Constructing the Wavelet Tree and Wavelet Matrix in Distributed Memory

The wavelet tree (Grossi et al. [SODA, 2003]) is a compact index for texts that provides rank, select, and access operations. This leads to many applications in text indexing, computational geometry, and compression. We present the first distributed memory wavelet tree construction algorithms, which allow us to process inputs that are orders of magnitude larger than what current shared memory construction algorithms can work with. In addition, our algorithms can easily be adapted to compute the wavelet matrix (Claude et al. [Inf. Syst., 47:15–32, 2015]), an alternative representation of the wavelet tree. In practice, one of our distributed memory wavelet matrix construction algorithms is the first parallel algorithm that can compute the wavelet matrix for alphabets of arbitrary size.

Patrick Dinklage
TU Dortmund University
Dortmund, Germany
patrick.dinklage@tu-dortmund.de

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

CP13

RecSplit: Minimal Perfect Hashing Via Recursive Splitting

A *minimal perfect hash function* bijectively maps a key set S out of a universe U into the first $|S|$ natural numbers. Minimal perfect hash functions are used, for example, to map irregularly-shaped keys, such as strings, in a compact space so that metadata can then be simply stored in an array. While it is known that just 1.44 bits per key are necessary to store a minimal perfect hash function, no published technique can go below 2 bits per key in practice. We propose a new technique for storing minimal perfect hash functions with expected linear construction time and expected constant lookup time that makes it possible to build for the first time, for example, structures which need 1.56 bits per key, that is, within 8.3% of the lower bound, in less than 2 ms per key. We show that instances of our construction are able to *simultaneously* beat the construction time, space usage and lookup time of the state-of-the-art data structure reaching 2 bits per key. Moreover, we provide parameter choices giving structures which are competitive with alternative, larger-size data structures in terms of space and lookup time. The construction of our data structures can be easily parallelized or mapped on distributed computational units (e.g., within the MapReduce framework), and structures larger than the available RAM can be directly built in mass storage.

Sebastiano Vigna, Emmanuel Esposito
Università degli Studi di Milano
sebastiano.vigna@unimi.it, emmanuel.espo@gmail.com

Thomas Mueller Graf
Independent
thomas.tom.mueller@gmail.com

CP13

Cost-Optimal Assignment of Elements in Genome-Scale Multi-Way Bucketed Cuckoo Hash Tables

We present the first practical algorithm to solve the minimum cost assignment problem for multi-way bucketed Cuckoo hashing with $h \geq 2$ hash functions and buckets that store $b \geq 1$ elements each. We minimize the average lookup cost over all stored elements, assuming that an element in the bucket indicated by its j-th hash function incurs a lookup cost of j cache misses. Our method is based on a combination of the Bellman-Ford and Hopcroft-Karp algorithms for finding minimum cost paths in the Cuckoo assignment graph, using multiple sources in parallel. We find a cost-optimal assignment for the 2.38 billion canonical DNA 25-mers of the human genome in 4 to 48 hours of CPU time, using up to 48 GB of RAM, at hash table loads between 50% and 99.9% for bucket sizes 3 to 8. For bucket size $b = 4$ and 95% load, we obtain optimal costs of ≈ 1.2 cache misses per stored element and 1.4 per element that is not present, using 7 hours of CPU time.

Jens Zentgraf
TU Dortmund
jens.zentgraf@tu-dortmund.de

Henning Timm, Sven Rahmann
University of Duisburg-Essen

henning.timm@uni-due.de, Sven.Rahmann@uni-due.de

CP14

Euclidean Bottleneck Bounded-Degree Spanning Tree Ratios

Inspired by the seminal works of Khuller et al. (STOC 1994) and Chan (SoCG 2003) we study the bottleneck version of the Euclidean bounded-degree spanning tree problem. A bottleneck spanning tree is a spanning tree whose largest edge-length is minimum, and a bottleneck degree-K spanning tree is a degree-K spanning tree whose largest edge-length is minimum. Let β_K be the supremum ratio of the largest edge-length of the bottleneck degree-K spanning tree to the largest edge-length of the bottleneck spanning tree, over all finite point sets in the Euclidean plane. It is known that $\beta_5 = 1$, and it is easy to verify that $\beta_2 \geq 2$, $\beta_3 \geq \sqrt{2}$, and $\beta_4 > 1.175$. It is implied by the Hamiltonicity of the cube of the bottleneck spanning tree that $\beta_2 \leq 3$. The degree-3 spanning tree algorithm of Ravi et al. (STOC 1993) implies that $\beta_3 \leq 2$. Andersen and Ras (Networks, 68(4):302314, 2016) showed that $\beta_4 \leq \sqrt{3}$. We present the following improved bounds: $\beta_2 \geq \sqrt{7}$, $\beta_3 \leq \sqrt{3}$, and $\beta_4 \leq \sqrt{2}$. As a result, we obtain better approximation algorithms for Euclidean bottleneck degree-3 and degree-4 spanning trees. As parts of our proofs of these bounds we present some structural properties of the Euclidean minimum spanning tree which are of independent interest.

Ahmad Biniaz
University of Windsor
ahmad.biniaz@gmail.com

CP14

Tightening Curves on Surfaces Monotonically with Applications

We prove the first polynomial bound on the number of *monotonic* homotopy moves required to tighten a collection of closed curves on any compact orientable surface, where the number of crossings in the curve is not allowed to increase at any time during the process. The best known upper bound before was exponential, which can be obtained by combining the algorithm of de Graaf and Schrijver [*JCTB*, 1997] together with an exponential upper bound on the number of possible surface maps. To obtain the new upper bound we apply tools from hyperbolic geometry, as well as operations in graph drawing algorithms—the cluster and pipe expansions—to the study of curves on surfaces. As corollaries, we present two efficient algorithms for curves and graphs on surfaces. First, we provide a polynomial-time algorithm to convert any given multicurve on a surface into minimal position. Such an algorithm only existed for single closed curves, and it is known that previous techniques do not generalize to the multicurve case. Second, we provide a polynomial-time algorithm to reduce any k-terminal plane graph using electrical transformations for arbitrary integer k. Previous algorithms only existed in the planar setting when $k \leq 4$, and all of them rely on extensive case-by-case analysis based on different values of k. Our algorithm makes use of the connection between electrical transformations and homotopy moves, and thus solves the problem in a unified fashion.

Hsien-Chih Chang
Department of Computer Science
Duke University
hcsoso@gmail.com

CP14

Embeddability of Simplicial Complexes Is Undecidable

We consider the following decision problem $\text{EMBED}_{k\to d}$ in computational topology (where $k \leq d$ are fixed positive integers): Given a finite simplicial complex K of dimension k, does there exist a (piecewise-linear) embedding of K into $\mathbb{R}^d$? The special case $\text{EMBED}_{1\to 2}$ is graph planarity, which is decidable in linear time, by the well-known algorithm of Hopcroft and Tarjan. In higher dimensions, $\text{EMBED}_{2\to 3}$ and $\text{EMBED}_{3\to 3}$ are known to be decidable (as well as NP-hard), and recent results of Čadek et al. imply that $\text{EMBED}_{k\to d}$ can be solved in polynomial time for any fixed pair (k, d) of dimensions in the so-called *metastable range* $d \geq \frac{3(k+1)}{2}$. Here, we prove that $\text{EMBED}_{k\to d}$ is algorithmically undecidable for almost all pairs of dimensions outside the metastable range, i.e. for $8 \leq d < \frac{3(k+1)}{2}$. This almost completely resolves the decidability vs. undecidability of $EMBED_{k\to d}$ in higher dimensions and establishes a dichotomy between polynomial-time solvability and undecidability. Our proof builds on work by Čadek et al., who showed how to encode an arbitrary system of Diophantine equations into a homotopy-theoretic *extension problem*; we turn their construction into an embeddability problem, using techniques from piecewise-linear (PL) topology due to Zeeman, Irwin, and others.

Marek Filakovsky
Masaryk University
Brno
mfilakov@ist.ac.at

Uli Wagner, Stephan Zhechev
IST Austria
uli@ist.ac.at, stephan.zhechev@ist.ac.at

CP14

Optimal Orthogonal Drawings of Planar 3-Graphs in Linear Time

A *planar orthogonal drawing* Γ of a planar graph G is a geometric representation of G such that the vertices are drawn as distinct points of the plane, the edges are drawn as chains of horizontal and vertical segments, and no two edges intersect except at their common end-points. A *bend* of Γ is a point of an edge where a horizontal and a vertical segment meet. Γ is *bend-minimum* if it has the minimum number of bends over all possible planar orthogonal drawings of G. This paper addresses a long standing, widely studied, open question: Given a planar 3-graph G (i.e., a planar graph with vertex degree at most three), what is the best computational upper bound to compute a bend-minimum planar orthogonal drawing of G in the variable embedding setting? In this setting the algorithm can choose among the exponentially many planar embeddings of G the one that leads to an orthogonal drawing with the minimum number of bends. We answer the question by describing an $O(n)$-time algorithm that computes a bend-minimum planar orthogonal drawing of G with at most one bend per edge, where n is the number of vertices of G. The existence of an orthogonal drawing algorithm that simultaneously minimizes the total number of bends and the number of bends per edge was previously unknown.

Walter Didimo, Giuseppe Liotta, Giacomo Ortali

Università degli Studi di Perugia
Italy
walter.didimo@unipg.it, giuseppe.liotta@unipg.it, giacomo.ortali@gmail.com

Maurizio Patrignani
Dipartimento di Informatica e Automazione,
Roma Tre University
maurizio.patrignani@uniroma3.it

CP14

Optimal Bound on the Combinatorial Complexity of Approximating Polytopes

Convex bodies play a fundamental role in geometric computation, and approximating such bodies is often a key ingredient in the design of efficient algorithms. We consider the question of how to succinctly approximate a multidimensional convex body by a polytope. We are given a convex body K of unit diameter in Euclidean d-dimensional space (where d is a constant) along with an error parameter $\varepsilon > 0$. The objective is to determine a polytope of low combinatorial complexity whose Hausdorff distance from K is at most ε. By *combinatorial complexity* we mean the total number of faces of all dimensions of the polytope. In the mid-1970's, a result by Dudley showed that $O(1/\varepsilon^{(d-1)/2})$ facets suffice, and Bronshteyn and Ivanov presented a similar bound on the number of vertices. While both results match known worst-case lower bounds, obtaining a similar upper bound on the total combinatorial complexity has been open for over 40 years. Recently, we made a first step forward towards this objective, obtaining a suboptimal bound. In this paper, we settle this problem with an asymptotically optimal bound of $O(1/\varepsilon^{(d-1)/2})$.

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

Rahul Arya
University of California, Berkeley
rahularya@berkeley.edu

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

Guilherme da Fonseca
Aix-Marseille Universite
guilherme.fonseca@lis-lab.fr

CP15

Near-Optimal Approximate Discrete and Continuous Submodular Function Minimization

In this paper we provide improved running times and oracle complexities for approximately minimizing a submodular function. Our main result is a randomized algorithm, which given any submodular function defined on n-elements with range $[-1, 1]$, computes an ε-additive approximate minimizer in $\tilde{O}(n/\varepsilon^2)$ oracle evaluations with high probability. This improves over the $\tilde{O}(n^{5/3}/\varepsilon^2)$ oracle evaluation algorithm of Chakrabarty et al. (STOC 2017) and the $\tilde{O}(n^{3/2}/\varepsilon^2)$ oracle evaluation algorithm of Hamoudi et al. Further, we leverage a generalization of this result to obtain efficient algorithms for minimizing a broad class of nonconvex functions. For any function f with domain $[0, 1]^n$ that satisfies $\frac{\partial^2 f}{\partial x_i \partial x_j} \leq 0$ for all $i \neq j$ and is L-Lipschitz with respect to the L^∞-norm we give an algorithm that computes an ε-additive approximate minimizer with $\tilde{O}(n \cdot \mathrm{poly}(L/\varepsilon))$ function evaluation with high probability.

Yang Liu, Aaron Sidford, Brian Axelrod
Stanford University
yangpatil@gmail.com, sidford@stanford.edu, baxelrod@stanford.edu

CP15

Sticky Brownian Rounding and Its Applications to Constraint Satisfaction Problems

Semidefinite programming is a powerful tool in the design and analysis of approximation algorithms for combinatorial optimization problems. In particular, the random hyperplane rounding method of Goemans and Williamson has been extensively studied for more than two decades, resulting in various extensions for a wide range of applications. Despite the fact that this approach yields tight approximation guarantees for some problems, e.g., Max-Cut, for many others, e.g., Max-SAT and Max-DiCut, the tight approximation ratio is still unknown. One of the main reasons for this is the fact that very few techniques for rounding semi-definite relaxations are known. In this work, we present a new general and simple method for rounding semi-definite programs, based on Brownian motion. Our approach is inspired by recent results in algorithmic discrepancy theory. We develop and present tools for analyzing our new rounding algorithms, utilizing mathematical machinery from the theory of Brownian motion, complex analysis, and partial differential equations. We apply our method to several classical problems, and derive new algorithms that are competitive with the best known results. To illustrate the versatility and general applicability of our approach, we give new approximation algorithms for the Max-Cut problem with side constraints that crucially utilizes measure concentration results for the Sticky Brownian Motion, a feature missing from hyperplane rounding and its generalizations.

Sepehr Abbasi-Zadeh
University of Toronto
sabbasizadeh@gmail.com

Nikhil Bansal
CWI
bansal@gmail.com

Guru Guruganesh
Google Research
gurug@google.com

Aleksandar Nikolov
University of Toronto
anikolov@cs.toronto.edu

Roy Schwartz
Technion
schwartz@cs.technion.ac.il

Mohit Singh
H. Milton Stewart School of Industrial & Systems Engineering
Georgia Institute of Technology

mohitsinghr@gmail.com

CP15

Spherical Discrepancy Minimization and Algorithmic Lower Bounds for Covering the Sphere

Inspired by the boolean discrepancy problem, we study the following optimization problem which we term Spherical Discrepancy: given m unit vectors $v_1, \dots, v_m$, find another unit vector x that minimizes $\max_i \langle x, v_i \rangle$. We show that Spherical Discrepancy is APX-hard and develop a multiplicative weights-based algorithm that achieves nearly optimal worst-case error bounds. We use our algorithm to give the first non-trivial lower bounds for the problem of covering a hypersphere by hyperspherical caps of uniform volume at least $2^{-o(\sqrt{n})}$. Up to a log factor, our lower bounds match known upper bounds in this regime.

Chris Jones, Matt McPartlon
University of Chicago
csj@uchicago.edu, mmcpartlon@uchicago.edu

CP15

Regular Languages Meet Prefix Sorting

Indexing strings via prefix (or suffix) sorting is, arguably, one of the most successful algorithmic techniques developed in the last decades. Can indexing be extended to languages? The main contribution of this paper is to initiate the study of the sub-class of regular languages accepted by an automaton whose states can be prefix-sorted. Starting from the recent notion of Wheeler graph - which extends naturally the concept of prefix sorting to labeled graphs - we investigate the properties of Wheeler languages, that is, regular languages admitting an accepting Wheeler finite automaton. We first characterize this family as the natural extension of regular languages endowed with the co-lexicographic ordering: the sorted prefixes of strings belonging to a Wheeler language are partitioned into a finite number of co-lexicographic intervals, each formed by elements from a single Myhill-Nerode equivalence class. We proceed by proving several results related to Wheeler automata, including a determinization algorithm for Wheeler NFAs which only doubles the number of states, polynomial-time sorting and minimization algorithms for Wheeler DFAs, and an algorithm that converts any acyclic DFA into the smallest equivalent Wheeler DFA. The latter contribution represents a provably minimum-size solution for the well-studied problem of indexing deterministic-acyclic graphs for linear-time pattern matching queries.

Nicola Prezza
University of Pisa
nicola.prezza@gmail.com

Jarno Alanko
University of Helsinki
jarno.alanko@helsinki.fi

Giovanna D'Agostino, Alberto Policriti
University of Udine
giovanna.dagostino@uniud.it, alberto.policriti@uniud.it

CP15

Faster p-Norm Minimizing Flows, Via Smoothed q-Norm Problems

We present faster high-accuracy algorithms for computing ℓ_p-norm minimizing flows. On a graph with m edges, our algorithm can compute a $(1+1/\text{poly}(m))$-approximate unweighted ℓ_p-norm minimizing flow with $pm^{1+\frac{1}{p-1}+o(1)}$ operations, for any $p \geq 2$, giving the best bound for all $p5.24$. Combined with the algorithm from the work of Adil et al. (SODA '19), we can now compute such flows for any $2 \leq p \leq m^{o(1)}$ in time at most $O(m^{1.24})$. In comparison, the previous best running time was $\Omega(m^{1.33})$ for large constant p. For $p \sim \delta^{-1} \log m$, our algorithm computes a $(1+\delta)$-approximate maximum flow on undirected graphs using $m^{1+o(1)}\delta^{-1}$ operations, matching the current best bound, albeit only for unit-capacity graphs. We also give an algorithm for solving general ℓ_p-norm regression problems for large p. Our algorithm makes $pm^{\frac{1}{3}+o(1)} \log^2(1/\varepsilon)$ calls to a linear solver. This gives the first high-accuracy algorithm for computing weighted ℓ_p-norm minimizing flows that runs in time $o(m^{1.5})$ for some $p = m^{\Omega(1)}$. Our key technical contribution is to show that smoothed ℓ_p-norm problems introduced by Adil et al., are interreducible for different values of p. No such reduction is known for standard ℓ_p-norm problems.

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

CP16

Detecting Feedback Vertex Sets of Size k in $O^\star(2.7^K)$ Time

In the Feedback Vertex Set problem, one is given an undirected graph G and an integer k, and one needs to determine whether there exists a set of k vertices that intersects all cycles of G (a so-called feedback vertex set). Feedback Vertex Set is one of the most central problems in parameterized complexity: It served as an excellent test bed for many important algorithmic techniques in the field such as Iterative Compression [Guo et al. (JCSS'06)], Randomized Branching [Becker et al. (J. Artif. Intell. Res'00)] and Cut&Count [Cygan et al. (FOCS'11)]. In particular, there has been a long race for the smallest dependence $f(k)$ in run times of the type $O^\star(f(k))$, where the $O^\star$ notation omits factors polynomial in n. This race seemed to be run in 2011, when a randomized $O^\star(3^k)$ time algorithm based on Cut&Count was introduced. In this work, we show the contrary and give a $O^\star(2.7^k)$ time randomized algorithm. Our algorithm combines all mentioned techniques with substantial new ideas: First, we show that, given a feedback vertex set of size k of bounded average degree, a tree decomposition of width $(1-\Omega(1))k$ can be found in polynomial time. Second, we give a randomized branching strategy inspired by the one from [Becker et al. (J. Artif. Intell. Res'00)] to reduce to the aforementioned bounded average degree setting. Third, we obtain significant run time improvements by employing fast matrix multiplication.

Jason M. Li
Carnegie Mellon University
jmli@cs.cmu.edu

Jesper Nederlof
TU Eindhoven
j.nederlof@tue.nl

CP16

A Nearly 5/3-Approximation Fpt Algorithm for

Min-K-Cut

Given an edged-weighted graph GGG, the min-kkk-cut problem asks for a set of edges with minimum total weight whose removal breaks the graph GGG into at least kkk connected components. It is well-known that the greedy algorithm can find a (2-2/k)(2-2/k)(2-2/k)-approximation of the min-kkk-cut in polynomial time. Assuming the Small Set Expansion Hypothesis (SSEH), no polynomial time algorithm can achieve an approximation ratio better than two [**?**]. Recently, Gupta, Lee and Li [**?**] gave a 1.99971.99971.9997-approximation FPT algorithm for the min-kkk-cut parameterized by kkk. They also improved this approximation ratio to 1.81 [**?**]. We generalize their proof techniques and show that the min-kkk-cut has a nearly 5/35/35/3-approximation FPT algorithm. Our proof is self-contained and much shorter than that of Gupta, Lee and Li.

Ken-ichi Kawarabayashi
National Institute of Informatics
JST, ERATO, Kawarabayashi Project
k_keniti@nii.ac.jp

Bingkai Lin
University of Tokyo
lin@nju.edu.cn

CP16

A $4+\epsilon$ Approximation for k-Connected Subgraphs

We obtain approximation ratio $4+\frac{2}{\ell} \approx 4+\frac{4\lg k}{\lg n - \lg k}$ for the (undirected) k-Connected Subgraph problem problem, where $\ell = \left\lfloor \frac{\lg n - \lg k + 1}{2\lg k + 1} \right\rfloor$ is the largest integer such that $2^{\ell-1}k^{2\ell+1} \leq n$. For large values of n this improves the ratio 6 of Cheriyan and Végh when $n \geq k^3$ (the case $\ell = 1$). Our result implies an fpt-approximation ratio $4+\epsilon$ that matches (up to the "$+\epsilon$' term) the best known ratio 4 for $k = 6, 7$ for both the general and the easier augmentation versions of the problem. Similar results are shown for the problem of covering an arbitrary crossing supermodular biset function.

Zeev Nutov
The Open University of Israel, Math and CS Dept.
.
nutov@openu.ac.il

CP16

Hitting Topological Minor Models in Planar Graphs is Fixed Parameter Tractable

For a finite collection of graphs $\mathcal{F}$, the $\mathcal{F}$-TM-Deletion problem has as input an n-vertex graph G and an integer k and asks whether there exists a set $S \subseteq V(G)$ with $|S| \leq k$ such that $G \setminus S$ does not contain any of the graphs in $\mathcal{F}$ as a topological minor. We prove that for every such $\mathcal{F}$, $\mathcal{F}$-TM-Deletion is fixed parameter tractable on planar graphs. In particular, we provide an $f(h,k) \cdot n^2$ algorithm where h is an upper bound to the vertices of the graphs in $\mathcal{F}$.

Petr Golovach
University of Bergen, Norway
petr.golovach@uib.no

Dimitrios Thilikos
CNRS, LIRMM, Montpellier
sedthilk@thilikos.info

Giannos Stamoulis
NKUA, Athens
giannosstam@di.uoa.gr

CP16

A Complexity Dichotomy for Hitting Connected Minors on Bounded Treewidth Graphs: the Chair and the Banner Draw the Boundary

For a fixed connected graph H, the $\{H\}$-M-Deletion problem asks, given a graph G, for the minimum number of vertices that intersect all minor models of H in G. It is known that this problem can be solved in time $f(\mathbf{tw}) \cdot n^{\mathcal{O}(1)}$, where $\mathbf{tw}$ is the treewidth of G. We determine the asymptotically optimal function $f(\mathbf{tw})$, for each possible choice of H. Namely, we prove that, under the ETH, $f(\mathbf{tw}) = 2^{\Theta(\mathbf{tw})}$ if H is a contraction of the chair or the banner, and $f(\mathbf{tw}) = 2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})}$ otherwise. Prior to this work, such a complete characterization was only known when H is a planar graph with at most five vertices. For the upper bounds, we present an algorithm in time $2^{\Theta(\mathbf{tw} \cdot \log \mathbf{tw})} \cdot n^{\mathcal{O}(1)}$ for the more general problem where all minor models of connected graphs in a finite family $\mathcal{F}$ need to be hit. We combine several ingredients such as the machinery of boundaried graphs in dynamic programming via representatives, the Flat Wall Theorem, Bidimensionality, the irrelevant vertex technique, treewidth modulators, and protrusion replacement. In particular, this algorithm vastly generalizes a result of Jansen et al. [SODA 2014] for the particular case $\mathcal{F} = \{K_5, K_{3,3}\}$.

Julien Baste
Institute of Optimization and Operations Research
Ulm University
julien.baste@uni-ulm.de

Ignasi Sau, Dimitrios Thilikos
CNRS, LIRMM, Montpellier
ignasi.sau@lirmm.fr, sedthilk@thilikos.info

CP17

One (more) Line on the Most Ancient Algorithm in History

We give a new simple and short ("one-line') analysis for the runtime of the well-known Euclidean Algorithm. While very short simple, the obtained upper bound in near-optimal.

Ilya Volkovich
Department of EECS, CSE Division
University of Michigan
ilyavol@umich.edu

Bruno Grenet
LIRMM, Université de Montpellier, CNRS
bruno.grenet@lirmm.fr

Vishwas Bhargava
Rutgers University

vishwas1384@gmail.com

CP17

Reducing 3SUM to Convolution-3SUM

Given a set S of n numbers, the 3SUM problem asks to determine whether there exist three elements $a, b, c \in S$ such that $a+b+c=0$. The related Convolution-3SUM problem asks to determine whether there exist a pair of indices i, j such that $A[i] + A[j] = A[i+j]$, where A is a given array of n numbers. When the numbers are integers, a *randomized* reduction from 3SUM to Convolution-3SUM was given in a seminal paper by Pătraşcu [STOC 2010], which was later improved by Kopelowitz, Pettie, and Porat [SODA 2016] with an $O(\log n)$ factor slowdown. In this paper, we present a simple *deterministic* reduction from 3SUM to Convolution-3SUM for integers bounded by U. We also describe additional ideas to obtaining further improved reductions, with only a $(\log\log n)^{O(1)}$ factor slowdown in the randomized case, and a $\log^{O(1)} U$ factor slowdown in the deterministic case.

Timothy M. Chan
University of Illinois at Urbana-Champaign
tmc@illinois.edu

Qizheng He
Department of Computer Science
University of Illinois at Urbana-Champaign
qizheng6@illinois.edu

CP17

Multiplicative Rank-1 Approximation Using Length-Squared Sampling

We show that the span of $\Omega(\frac{1}{\varepsilon^4})$ rows of any matrix $A \subset \mathbb{R}^{n\times d}$ sampled according to the length-squared distribution contains a rank-1 matrix $\tilde{A}$ such that $||A - \tilde{A}||_F^2 \leq (1+\varepsilon) \cdot ||A - \pi_1(A)||_F^2$, where $\pi_1(A)$ denotes the best rank-1 approximation of A under the Frobenius norm. Length-squared sampling has previously been used in the context of rank-k approximation. However, the approximation obtained was additive in nature. We obtain a multiplicative approximation albeit only for rank-1 approximation.

Ragesh Jaiswal
Indian Institute of Technology Delhi
rjaiswal@gmail.com

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

CP17

Quantum Approximate Counting, Simplified

In 1998, Brassard, Høyer, Mosca, and Tapp (BHMT) gave a quantum algorithm for *approximate counting*. Given a list of N items, K of them marked, their algorithm estimates K to within relative error ε by making only $O\left(\frac{1}{\varepsilon}\sqrt{\frac{N}{K}}\right)$ queries. Although this speedup is of "Grover" type, the BHMT algorithm has the curious feature of relying on the Quantum Fourier Transform (QFT), more commonly associated with Shor's algorithm. Is this necessary? This paper presents a simplified algorithm, which we prove achieves the same query complexity using Grover iterations only. We also generalize this to a QFT-free algorithm for amplitude estimation. Related approaches to approximate counting were sketched previously by Grover, Abrams and Williams, Suzuki et al., and Wie (the latter two as we were writing this paper), but in all cases without rigorous analysis.

Scott Aaronson, Patrick J. Rall
University of Texas at Austin
aaronson@cs.utexas.edu, patrickjrall@gmail.com

CP17

Bucket Oblivious Sort: An Extremely Simple Oblivious Sort

We propose a conceptually simple oblivious sort and oblivious random permutation algorithms called bucket oblivious sort and bucket oblivious random permutation. Bucket oblivious sort uses 6n log n time (measured by the number of memory accesses) and 2Z client storage to ensure an error probability exponentially small in Z. This runtime is only 3x slower than a non-oblivious merge sort baseline; for 2^{30} elements, it is 5x faster than bitonic sort, the de facto oblivious sorting algorithm in practical implementations.

Ling Ren
University of Illinois Urbana-Champaign
renling@illinois.edu

CP18

Lossless Prioritized Embeddings

Given metrics (X, d) and (Y, ρ) and an ordering $x_1, x_2, \ldots, x_n$ of X, an embedding $f : X \to Y$ is said to have a *prioritized distortion* $\alpha(\cdot)$, for a function $\alpha()$, if for any pair x_j, x' of distinct points in X, the distortion provided by f for this pair is at most $\alpha(j)$. In EFN15 a general methodology for constructing prioritized embeddings was developed. Though this methodology enables one to come up with many prioritized embeddings, it typically incurs some loss in the distortion. This loss is problematic for isometric embeddings. It is also troublesome for Matousek's embedding of general metrics into ℓ_∞, which for a parameter $k = 1, 2, \ldots$, provides distortion $2k-1$ and dimension $O(k \log n \cdot n^{1/k})$. In this paper we devise two *lossless* prioritized embeddings. The first one is an isometric prioritized embedding of tree metrics into ℓ_∞ with dimension $O(\log j)$, matching the worst-case guarantee of the classical embedding of LLR95. The second one is a prioritized Matousek's embedding of general metrics into ℓ_∞, which for a parameter $k = 1, 2, \ldots$, provides prioritized distortion $2\lceil k\frac{\log j}{\log n}\rceil - 1$ and dimension $O(k \log n \cdot n^{1/k})$, again matching the worst-case guarantee of the classical Matousek's embedding.

Michael Elkin
Ben-Gurion University
elkinm@cs.bgu.ac.il

CP18

Labelings Vs. Embeddings: On Distributed Representations of Distances

We investigate for which metric spaces the performance of distance labeling and of ℓ_∞-embeddings differ, and how significant can this difference be. A distance labeling is a distributed representation of distances in a metric space (X, d), where each point $x \in X$ is assigned a succinct label, such that the distance between any two points $x, y \in X$

can be approximated given only their labels. A highly structured special case is an embedding into ℓ_∞, where each point $x \in X$ is assigned a vector $f(x)$ such that $\|f(x)-f(y)\|_\infty$ is approximately $d(x,y)$. The performance of a distance labeling or an ℓ_∞-embedding is measured via its distortion and its label-size/dimension. We also study the analogous question for the prioritized versions of these two measures. Here, a priority order $\pi = (x_1, \ldots, x_n)$ of the point set X is given, and higher-priority points should have shorter labels. Finally, we compare prioritized versions to their classical ones. We answer these questions in several scenarios, uncovering a surprisingly diverse range of behaviors. First, in some cases labelings and embeddings have very similar worst-case performance, but in other cases there is a huge disparity. In the prioritized setting, we most often find a strict separation. And finally, when comparing the classical and prioritized settings, we find that the worst-case bound for label size often "translates' to a prioritized one.

Arnold Filtser
Columbia University
arnold273@gmail.com

Lee-Ad Gottlieb
Ariel University
leead@ariel.ac.il

Robert Krauthgamer
Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il

CP18

Quasi-Polynomial Algorithms for Submodular Tree Orienteering and Other Directed Network Design Problems

We consider the following general network design problem on directed graphs. The input is an asymmetric metric (V, c), root $r^* \in V$, monotone submodular function $f : 2^V \to \mathbb{R}_+$ and budget B. The goal is to find an r^*-rooted arborescence T of cost at most B that maximizes $f(T)$. Our main result is a very simple quasi-polynomial time $O(\frac{\log k}{\log\log k})$-approximation algorithm for this problem, where $k \leq |V|$ is the number of vertices in an optimal solution. To the best of our knowledge, this is the first non-trivial approximation ratio for this problem. As a consequence we obtain an $O(\frac{\log^2 k}{\log\log k})$-approximation algorithm for directed (polymatroid) Steiner tree in quasi-polynomial time. We also extend our main result to a setting with additional length bounds at vertices, which leads to improved $O(\frac{\log^2 k}{\log\log k})$-approximation algorithms for the single-source buy-at-bulk and priority Steiner tree problems. For the usual directed Steiner tree problem, our result matches the best previous approximation ratio [GLL18], but improves significantly on the running time: our algorithm takes $n^{O(\log^{1+\epsilon} k)}$ time whereas the previous algorithm required $n^{O(\log^5 k)}$ time. Under certain complexity assumptions, our approximation ratios are best possible (up to constant factors).

Rohan Ghuge, Viswanath Nagarajan
University of Michigan
rghuge@umich.edu, viswa@umich.edu

CP18

2-Approximating Feedback Vertex Set in Tournaments

A *tournament* is a directed graph T such that every pair of vertices is connected by an arc. A *feedback vertex set* is a set S of vertices in T such that $T - S$ is acyclic. We consider the FEEDBACK VERTEX SET problem in tournaments. Here the input is a tournament T and a weight function $w : V(T) \to \mathbb{N}$ and the task is to find a feedback vertex set S in T minimizing $w(S) = \sum_{v \in S} w(v)$. Rounding optimal solutions to the natural LP-relaxation of this problem yields a simple 3-approximation algorithm. This has been improved to 2.5 by Cai et al. [SICOMP 2000], and subsequently to 7/3 by Mnich et al. [ESA 2016]. In this paper we give the first polynomial time factor 2 approximation algorithm for this problem. Assuming the Unique Games conjecture, this is the best possible approximation ratio achievable in polynomial time.

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

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Pranabendu Misra
Max Planck Institute for Informatics
pranabendu.m@gmail.com

Joydeep Mukherjee
ISI
joydeep.m1981@gmail.com

Fahad Panolan
IIT Hyderabad
fahad.panolan@gmail.com

Geevarghese Philip
Chennai Mathematical Institute
Chennai, India
gphilip@cmi.ac.in

CP18

Fast Lp-Based Approximations for Geometric Packing and Covering Problems

We derive fast approximation schemes for LP relaxations of several well-studied geometric optimization problems that include packing, covering, and mixed packing and covering constraints. Previous work in computational geometry concentrated mainly on the rounding stage to proveapproximation bounds, assuming that the underlying LPs can be solved efficiently. This work demonstrates that many of those results can be made to run in nearly linear time. In contrast to prior work on this topic our algorithms handle weights and capacities, side constraints, and also apply to mixed packing and covering problems, in a unified fashion. Our framework relies crucially on the properties of a randomized MWU algorithm of [CQ]; we demonstrate that it is well-suited for range spaces that admit efficient approximate dynamic data structures foremptiness oracles. Our framework cleanly separates the MWU algorithm for solving the LP from the key geometric data structure primitives, and this enables us to handle side constraints in a simple way. Combined with rounding algorithms that can also be implemented efficiently, we obtain the first near-linear time constant factor approximation algorithms for several problems.

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

Sariel Har-Peled
Department of Computer Science
University of Illinois, Urbana Champaign
sariel@illinois.edu

CP19

Round Complexity of Common Randomness Generation: The Amortized Setting

We study the effect of rounds of interaction on the common randomness generation (CRG) problem. In the CRG problem, two parties receive samples X_i and Y_i, respectively, drawn jointly from a distribution μ. The two parties wish to agree on a common random key of high entropy by exchanging messages over multiple rounds. In this work we study the amortized version of the problem, i.e., the amount of communication required to generate the random key in the limit as the number of bits generated tends to infinity. Recently Bafna et al. (SODA 2019) considered the non-amortized version of the problem: they gave, for $r, n \in \mathbf{N}$, a source $\mu_{r,n}$, such that with $r+2$ rounds one can generate n bits of common randomness from $\mu_{r,n}$ with $O(r \log n)$ communication, whereas with roughly $r/2$ rounds the communication complexity is $\Omega(n/\text{poly} \log n)$. In this work we strengthen the work of Bafna et al. in two ways: First we show that the results extend to the amortized setting. We also reduce the gap between the round complexity in the upper and lower bounds to an additive constant. Specifically we show that for every pair $r, n \in \mathbf{N}$ the amortized communication complexity to generate $\Omega(n)$ bits of common randomness from the source $\mu_{r,n}$ using $r+2$ rounds of communication is $O(r \log n)$ whereas the amortized communication required to generate the same amount of randomness from r rounds is $\Omega(\sqrt{n})$.

Noah Golowich
Harvard University
nzg@mit.edu

Madhu Sudan
Harvard
madhu@cs.harvard.edu

CP19

A New Lower Bound on Hadwiger-Debrunner Numbers in the Plane

A family of sets F is said to satisfy the (p, q) property if among any p sets in F some q have a non-empty intersection. Hadwiger and Debrunner (1957) conjectured that for any $p \geq q \geq d+1$ there exists $c = c_d(p, q)$, such that any family of compact convex sets in $\mathbb{R}^d$ that satisfies the (p, q) property can be pierced by at most c points. In a celebrated result from 1992, Alon and Kleitman proved the conjecture. However, obtaining sharp bounds on $c_d(p, q)$, known as the 'the HD numbers', is still a major open problem in discrete and computational geometry. The best currently known upper bound on the HD numbers in the plane is $O(p^{(1.5+\delta)(1+\frac{1}{q-2})})$, obtained by combination of a result of Keller, Smorodinsky, and Tardos and Rubin. The best lower bound is $c_2(p, q) = \Omega(\frac{p}{q} \log(\frac{p}{q}))$, obtained by Bukh, Matoušek and Nivasch. In this paper we improve the lower bound significantly by showing that $c_2(p, q) \geq p^{1+\Omega(\frac{1}{q})}$. Unlike previous bounds on the HD numbers, which mainly used the weak epsilon-net theorem, our bound stems from a surprising connection of the (p, q) problem to an old problem of Erdős on points in general position in the plane. We use a novel construction for Erdős' problem, obtained recently by Balogh and Solymosi using the hypergraph container method, to get the lower bound on $c_2(p, 3)$. We then generalize the bound to $c_2(p, q)$ for any $q \geq 3$.

Chaya Keller, Shakhar Smorodinsky
Ben Gurion University
chayak@ariel.ac.il, shakhar@math.bgu.ac.il

CP19

The Power of Distributed Verifiers in Interactive Proofs

We explore the power of interactive proofs with a distributed verifier. In this setting, the verifier consists of n nodes and a graph G that defines their communication pattern. The prover is a single entity that communicates with all nodes by short messages. The goal is to verify that the graph G belongs to some language in a small number of rounds, and with small communication bound, i.e., the proof size. This interactive model was introduced by Kol, Oshman and Saxena (PODC 2018) as a generalization of non-interactive distributed proofs. In this work, we provide a new general framework for distributed interactive proofs that allows one to translate standard interactive protocols (i.e., with a centralized verifier) to ones where the verifier is distributed with a proof size that depends on the computational complexity of the verification algorithm run by the centralized verifier. We show that every (centralized) computation implemented by either a small space or by uniform NC circuit can be translated into a distributed protocol with $O(1)$ rounds and $O(\log n)$ bits proof size for the low space case and $polylog(n)$ many rounds and proof size for NC. For many problems we show how to reduce proof size below the naturally seeming barrier of $\log n$. Our compilers capture many natural problems and demonstrates the difficultly in showing lower bounds in these regimes.

Eylon Yogev
Weizmann Institute
eylony@gmail.com

Merav Parter
The Weizmann Institute of Science, Rehovot, Israel.
merav.parter@weizmann.ac.il

Moni Naor
Weizmann Institute of Science
moni.naor@weizmann.ac.il

CP19

The Combinatorics of the Longest-Chain Rule: Linear Consistency for Proof-of-Stake Blockchains

The blockchain data structure maintained via the longest-chain rule–popularized by Bitcoin–is a powerful algorithmic tool for consensus algorithms. Such algorithms achieve consistency for blocks in the chain as a function of their depth from the end of the chain. While the analysis of Bitcoin guarantees consistency with error 2^{-k} for blocks of depth $O(k)$, the state-of-the-art of proof-of-stake (PoS) blockchains suffers from a quadratic dependence on k: these protocols, exemplified by Ouroboros (Crypto 2017), Ouroboros Praos (Eurocrypt 2018) and Sleepy Consensus

(Asiacrypt 2017), can only establish that depth $\Theta(k^2)$ is sufficient. Whether this quadratic gap is an intrinsic limitation of PoS–due to issues such as the nothing-at-stake problem–has been an urgent open question, as deployed PoS blockchains further rely on consistency for protocol correctnes. We give an axiomatic theory of blockchain dynamics that permits rigorous reasoning about the longest-chain rule and achieve, in broad generality, $\Theta(k)$ dependence on depth in order to achieve consistency error 2^{-k}. In particular, for the first time we show that PoS protocols can match proof-of-work protocols for linear consistency. We analyze the associated stochastic process, give a recursive relation for the critical functionals of this process, and derive tail bounds in both i.i.d. and martingale settings via associated generating functions.

Erica Blum
University of Maryland, College Park
erblum@cs.umd.edu

Aggelos Kiayias
IOHK
University of Edinburgh
akiayias@inf.ed.ac.uk

Cristopher Moore
Santa Fe Institute
cris@santafe.edu

Saad Quader, Alexander Russell
University of Connecticut
saad.quader@uconn.edu, acr@cse.uconn.edu

CP19

Lower Bounds for Oblivious Near-Neighbor Search

We prove an $\Omega(d \lg n/(\lg\lg n)^2)$ lower bound on the dynamic cell-probe complexity of statistically *oblivious* approximate-near-neighbor search (ANN) over the d-dimensional Hamming cube. For the natural setting of $d = \Theta(\log n)$, our result implies an $\tilde{\Omega}(\lg^2 n)$ lower bound, which is a quadratic improvement over the highest (non-oblivious) cell-probe lower bound for ANN. This is the first super-logarithmic *unconditional* lower bound for ANN against general (non black-box) data structures. We also show that any oblivious *static* data structure for decomposable search problems (like ANN) can be obliviously dynamized with $O(\log n)$ overhead in update and query time, strengthening a classic result of Bentley and Saxe (Algorithmica, 1980).

Kasper Green Larsen
Aarhus University
larsen@cs.au.dk

Tal Malkin, Omri Weinstein
Columbia University
tal@cs.columbia.edu, omri@cs.columbia.edu

Kevin Yeo
Google
kwlyeo@google.com

CP20

Combinatorial Generation Via Permutation Languages

We present a general and versatile algorithmic framework for exhaustively generating a large variety of different combinatorial objects, based on encoding them as permutations. This approach unifies many known results, including four classical Gray codes for permutations, bitstrings, binary trees and set partitions, and allows us to prove many new ones. We present two distinct new applications for our new framework: The first is the generation of pattern-avoiding permutations, yielding new Gray codes for different families of permutations characterized by the avoidance of certain classical patterns, (bi)vincular patterns, barred patterns, Bruhat-restricted patterns, mesh patterns, monotone and geometric grid classes etc. We thus obtain new Gray code algorithms for the combinatorial objects in bijection to these permutations, e.g., five different types of geometric rectangulations. The second main application are lattice congruences of the weak order on the symmetric group S_n. Recently, Pilaud and Santos realized all those lattice congruences as $(n-1)$-dimensional polytopes, called quotientopes, which generalize hypercubes, associahedra, permutahedra etc. Our algorithm generates the equivalence classes of each of those lattice congruences, by producing a Hamilton path on the skeleton of the corresponding quotientope. We thus obtain a provable notion of optimality for the Gray codes obtained from our framework: They translate into walks along the edges of a polytope.

Elizabeth Hartung
Massachusetts College of Liberal Arts
e.hartung@mcla.edu

Hung P. Hoang
Departmenr of Computer Science, ETH Zurich
hung.hoang@inf.ethz.ch

Torsten Mütze
Department of Computer Science, University of Warwick
torsten.mutze@warwick.ac.uk

Aaron Williams
Computer Science Department, Williams College
aw14@williams.edu

CP20

Navigating An Infinite Space with Unreliable Movements

We consider a search problem on a 2-dimensional infinite grid with a single mobile agent. The goal of the agent is to find her way home, which is located in a grid cell chosen by an adversary. Initially, the agent is provided with an infinite sequence of *instructions*, that dictate the movements performed by the agent. Each instruction corresponds to a movement to an adjacent grid cell and the set of instructions can be a function of the initial locations of the agent and home. The challenge of our problem stems from faults in the movements made by the agent. In every step, with some constant probability $0 \leq p \leq 1$, the agent performs a random movement instead of following the current instruction. This paper provides two results on this problem. First, we show that for some values of p, there does not exist any set of instructions that guide the agent home in finite expected time. Second, we complement this impossibility result with an algorithm that, for sufficiently small values of p, yields a finite expected hitting time for home. In particular, we show that for any $p < 1$, our approach gives a hitting rate that decays polynomially as a function of time. In that sense, our approach is far superior to a standard random walk in terms of hitting time. The main contribution and take-home message of this paper is to show that, for some value of $0.01139\dots < p < 0.6554\dots$, there exists a phase transition on the solvability of the problem.

Jara Uitto, Anders Martinsson
ETH Zurich
jara.uitto@aalto.fi, anders.martinsson@inf.ethz.ch

CP20

Edge Expansion and Spectral Gap of Nonnegative Matrices

The classic graphical Cheeger inequalities state that if M is an $n \times n$ symmetric doubly stochastic matrix, then

$$\frac{1-\lambda_2(M)}{2} \leq \phi(M) \leq \sqrt{2\cdot(1-\lambda_2(M))}$$

where $\phi(M) = \min_{S\subseteq[n],|S|\leq n/2}\left(\frac{1}{|S|}\sum_{i\in S, j\notin S} M_{i,j}\right)$ is the edge expansion of M, and $\lambda_2(M)$ is the second largest eigenvalue of M. We study the relationship between $\phi(A)$ and the spectral gap $1-\lambda_2(A)$ for any doubly stochastic matrix A (not necessarily symmetric), where $\lambda_2(A)$ is a nontrivial eigenvalue of A with maximum real part. Fiedler showed that the upper bound on $\phi(A)$ is unaffected, i.e., $\phi(A) \leq \sqrt{2\cdot(1-\lambda_2(A))}$, and for the lower bound on $\phi(A)$, there are known constructions with

$$\phi(A) \in \Theta\left(\frac{1-\lambda_2(A)}{\log n}\right).$$

In our first result, we provide an exponentially better construction of $n\times n$ doubly stochastic matrices A_n, for which

$$\phi(A_n) \leq \frac{1-\lambda_2(A_n)}{\sqrt{n}}.$$

We further show that for any doubly stochastic matrix A,

$$\phi(A) \geq \frac{1-\lambda_2(A)}{35\cdot n}.$$

Our second result extends these bounds to general nonnegative matrices R, relating the edge expansion (a quantitative measure of the irreducibility of R) and spectral gap of R, to obtain a two-sided quantitative refinement of the Perron-Frobenius theorem.

Jenish C. Mehta
California Institute of Technology
jenishc@gmail.com

Leonard Schulman
Caltech
schulman@caltech.edu

CP20

X-Ramanujan Graphs

Let X be an infinite graph of bounded degree; e.g., the Cayley graph of a free product of finite groups. If G is a finite graph covered by X, it is said to be X-Ramanujan if its second-largest eigenvalue $\lambda_2(G)$ is at most the spectral radius $\rho(X)$ of X. In case X is the infinite Δ-regular tree, this reduces to the well known notion of a finite Δ-regular graph being Ramanujan. Inspired by the Interlacing Polynomials method of Marcus, Spielman, and Srivastava, we show the existence of infinitely many X-Ramanujan graphs for a variety of infinite X. In particular, X need not be a tree; our analysis is applicable whenever X is what we call an additive product graph. This additive product is a new construction of an infinite graph $\mathsf{AddProd}(A_1, , A_c)$ from finite 'atom' graphs $A_1, , A_c$ over a common vertex set. It generalizes the notion of the free product graph $A_1 * \cdots * Ac$ when the atoms A_j are vertex-transitive, and it generalizes the notion of the universal covering tree. Key to our analysis is a new graph polynomial $\alpha(A_1,\dots,A_c;x)$ that we call the additive characteristic polynomial, which generalizes the well known matching polynomial $\mu(G;x)$. We show that all of its roots have magnitude at most $\rho(\mathsf{AddProd}(A_1, , A_c))$. This last fact is proven by generalizing Godsil's notion of treelike walks on a graph G to a notion of freelike walks on a collection of atoms $A_1,\dots,A_c$.

Sidhanth Mohanty
University of California Berkeley
sidhanthm@berkeley.edu

Ryan O'Donnell
Carnegie Mellon University, USA
odonnell@cs.cmu.edu

CP20

Linear Rankwidth Meets Stability

Classes with bounded rankwidth are MSO-transductions of trees and classes with bounded linear rankwidth are MSO-transductions of paths. These results show a strong link between the properties of these graph classes considered from the point of view of structural graph theory and from the point of view of finite model theory. We take both views on classes with bounded linear rankwidth and prove structural and model theoretic properties of these classes: 1) Graphs with linear rankwidth at most r are linearly χ-bounded. Actually, they have bounded c-chromatic number, meaning that they can be colored with $f(r)$ colors, each color inducing a cograph. 2) Based on a Ramsey-like argument, we prove for every proper hereditary family $\mathcal{F}$ of graphs (like cographs) that there is a class with bounded rankwidth that does not have the property that graphs in it can be colored by a bounded number of colors, each inducing a subgraph in $\mathcal{F}$. 3) For a class $\mathcal{C}$ with bounded linear rankwidth the following conditions are equivalent: a) $\mathcal{C}$ is stable, b) $\mathcal{C}$ excludes some half-graph as a semi-induced subgraph, c) $\mathcal{C}$ is a first-order transduction of a class with bounded pathwidth. These results open the perspective to study classes admitting low linear rankwidth covers.

Patrice Ossona de Mendez
Centre d'Analyse et de Mathematiques Sociales
CNRS
pom@ehess.fr

Jaroslav Nesetril
Institute for Theoretical Computer Science
Charles University Prague
nesetril@iuuk.mf.cuni.cz

Roman Rabinovich
Technical University Berlin
roman.rabinovich@tu-berlin.de

Sebastian Siebertz
University of Bremen
siebertz@uni-bremen.de

CP21

A Short Proof of the Toughness of Delaunay Tri-

angulations

We present a self-contained short proof of the seminal result of Dillencourt (SoCG 1987 and DCG 1990) that Delaunay triangulations, of planar point sets in general position, are 1-tough. An important implication of this result is that Delaunay triangulations have perfect matchings. Another implication of our result is a proof of the conjecture of Aichholzer et al. (2010) that at least n points are required to block any n-vertex Delaunay triangulation.

Ahmad Biniaz
University of Windsor
ahmad.biniaz@gmail.com

CP21

Dynamic Generalized Closest Pair: Revisiting Eppstein's Technique

Eppstein (1995) gave a technique to transform any data structure for dynamic nearest neighbor queries into a data structure for dynamic closest pair, for any distance function; the transformation increases the time bound by two logarithmic factors. We present a similar, simple transformation that is just as good, and can avoid the extra logarithmic factors when the query and update time of the given structure exceed n^{ε} for some constant $\varepsilon > 0$. Consequently, in the case of an arbitrary distance function, we obtain an optimal $O(n)$-space data structure to maintain the dynamic closest pair of n points in $O(n)$ amortized time plus $O(n)$ distance evaluations per update.

Timothy M. Chan
University of Illinois at Urbana-Champaign
tmc@illinois.edu

CP21

On the Change-Making Problem

Given a set of n non-negative integers representing a coin system, the *change-making problem* seeks the fewest number of coins that sum up to a given value t, where each type of coin can be used an unlimited number of times. This problem is a popular homework exercise in dynamic programming, where the textbook solution runs in $O(nt)$ time. It is not hard to solve this problem in $O(t \operatorname{polylog} t)$ time by using convolution. In this paper, we present a simple deterministic $O(t \log t \log \log t)$ time algorithm, and later improve the running time to $O(t \log t)$ by randomization.

Timothy M. Chan
University of Illinois at Urbana-Champaign
tmc@illinois.edu

Qizheng He
Department of Computer Science
University of Illinois at Urbana-Champaign
qizheng6@illinois.edu

CP21

Adaptive Discrete Phase Retrieval

In the phase retrieval problem, the goal is to recover an unknown signal vector $x \in C^N$ from a small number of measurements $\{y_i\}$ of the form $y_i = |\langle m_i, x\rangle|^2$, where $m_i \in C^N$ are measurement vectors. We introduce two variations of the traditional model: the *adaptive* setting where measurement vectors can depend on previous measurements, and the *discrete* setting where each component of x is representable using a bounded number of bits. In contrast to the heavy machinery used in prior work on phase retrieval, we design simple ensembles of measurement vectors (both adaptive and deterministic) for discrete phase retrieval. The number of samples needed is significantly lower than traditional phase retrieval. Our results highlight the role of bit precision in reasoning about the sample complexity of the phase retrieval problem.

Moses Charikar
Stanford University, California
moses@cs.stanford.edu

Xian Wu, Yinyu Ye
Stanford University
xwu20@stanford.edu, yyye@stanford.edu

CP21

Fast Fourier Sparsity Testing

A function $f : F_2^n \to R$ is s-*sparse* if it has at most s non-zero Fourier coefficients. Motivated by applications to fast sparse Fourier transforms over F_2^n, we study efficient algorithms for the problem of approximating the ℓ_2-distance from a given function to the closest s-sparse function. While previous works (e.g., Gopalan *et al.* SICOMP 2011) study the problem of distinguishing s-sparse functions from those that are far from s-sparse under Hamming distance, to the best of our knowledge no prior work has explicitly focused on the more general problem of distance estimation in the ℓ_2 setting, which is particularly well-motivated for noisy Fourier spectra. Given the focus on efficiency, our main result is an algorithm that solves this problem with query complexity $O(s)$ for constant accuracy and error parameters, which is only quadratically worse than applicable lower bounds.

Samson Zhou
Carnegie Mellon University
samsonzhou@gmail.com

Grigory Yaroslavtsev
University of Pennsylvania
grigory@grigory.us

CP22

Faster Deterministic Distributed Coloring Through Recursive List Coloring

We provide novel deterministic distributed vertex coloring algorithms. As our main result, we give a deterministic distributed algorithm to compute a $(\Delta + 1)$-coloring of an n-node graph with maximum degree Δ in $2^{O(\sqrt{\log \Delta})} \cdot \log n$ rounds. For graphs with arboricity a, we obtain a deterministic distributed algorithm to compute a $(2 + o(1))a$-coloring in time $2^{O(\sqrt{\log a})} \cdot \log^2 n$. Further, for graphs with bounded neighborhood independence, we show that a $(\Delta + 1)$-coloring can be computed more efficiently in time $2^{O(\sqrt{\log \Delta})} + O(\log^* n)$. This in particular implies that also a $(2\Delta - 1)$-edge coloring can be computed deterministically in $2^{O(\sqrt{\log \Delta})} + O(\log^* n)$ rounds, which improves the best known time bound for small values of Δ. All results even hold for the list coloring variants of the problems. As a consequence, we also obtain an improved deterministic $2^{O(\sqrt{\log \Delta})} \cdot \log^3 n$-round algorithm for Δ-coloring non-complete graphs with maximum degree $\Delta \geq 3$. Most of our algorithms only require messages of $O(\log n)$ bits (includ-

ing the $(\Delta+1)$-vertex coloring algorithms). Our main technical contribution is a recursive deterministic distributed list coloring algorithm to solve list coloring problems with lists of size $\Delta^{1+o(1)}$.

Fabian Kuhn
University of Freiburg, Germany
kuhn@cs.uni-freiburg.de

CP22

Faster Algorithms for Edge Connectivity Via Random 2-Out Contractions

We provide a simple new randomized contraction approach to the global minimum cut problem for simple undirected graphs. The contractions exploit 2-out edge sampling from each vertex rather than the standard uniform edge sampling. We demonstrate the power of our new approach by obtaining better algorithms for sequential, distributed, and parallel models of computation. Our end results include the following randomized algorithms for computing edge connectivity, with high probability:

- Two *sequential* algorithms with complexities $O(m \log n)$ and $O(m + n \log^3 n)$. These improve on a long line of developments including a celebrated $O(m \log^3 n)$ algorithm of Karger [STOC'96] and the state of the art $O(m \log^2 n(\log \log n)^2)$ algorithm of Henzinger et al. [SODA'17]. Moreover, our $O(m + n \log^3 n)$ algorithm is optimal when $m = \Omega(n \log^3 n)$.

- An $\tilde{O}(n^{0.8} D^{0.2} + n^{0.9})$ round *distributed* algorithm, where D denotes the graph diameter. This improves substantially on a recent breakthrough of Daga et al.[STOC'19], which achieved a round complexity of $\tilde{O}(n^{1-1/353} D^{1/353} + n^{1-1/706})$, hence providing the first sublinear distributed algorithm for exactly computing the edge connectivity.

- The first $O(1)$ round algorithm for the *massively parallel computation* setting with linear memory per machine.

Krzysztof Nowicki
University of Wroclaw
knowicki@cs.uni.wroc.pl

Mohsen Ghaffari
ETH Zurich
ghaffari@inf.ethz.ch

Mikkel Thorup
University of Copenhagen
mikkel2thorup@gmail.com

CP22

Finding a Bounded-Degree Expander Inside a Dense One

It follows from the Marcus-Spielman-Srivastava proof of the Kadison-Singer conjecture that if $G = (V, E)$ is a Δ-regular dense expander then there is an edge-induced subgraph $H = (V, E_H)$ of G of constant maximum degree which is also an expander. As with other consequences of the MSS theorem, it is not clear how one would explicitly construct such a subgraph. We show that such a subgraph (although with quantitatively weaker expansion and near-regularity properties than those predicted by MSS) can be constructed with high probability in linear time, via a simple algorithm. Our algorithm allows a distributed implementation that runs in $\mathcal{O}(\log n)$ rounds and does $\mathcal{O}(n)$ total work with high probability. The analysis of the algorithm is complicated by the complex dependencies that arise between edges and between choices made in different rounds. We sidestep these difficulties by following the combinatorial approach of counting the number of possible random choices of the algorithm which lead to failure. We do so by a compression argument showing that such random choices can be encoded with a non-trivial compression. Our algorithm bears some similarity to the way agents construct a communication graph in a peer-to-peer network, and, in the bipartite case, to the way agents select servers in blockchain protocols.

Francesco Pasquale
Università di Roma Tor Vergata
pasquale@mat.uniroma2.it

Luca Becchetti
Sapienza Università di Roma
becchetti@diag.uniroma2.it

Andrea Clementi
University of Rome Tor Vergata
clementi@mat.uniroma2.it

Emanuele Natale
Université Côte dAzur, CNRS, INRIA
natale@unice.fr

Luca Trevisan
Bocconi University
l.trevisan@unibocconi.it

CP22

Parallel Batch-Dynamic Graphs: Algorithms and Lower Bounds

We study the problem of dynamically maintaining graph properties under batches of edge insertions and deletions in the massively parallel model of computation. In this setting, the graph is stored on a number of machines, each having space strongly sublinear with respect to the number of vertices. Our goal is to handle batches of updates and queries in constant rounds of parallel computation, as well as to reduce the total communication between the machines. This objective corresponds to the gradual buildup of databases, while obtaining constant rounds of communication for problems in the static setting has been elusive for problems as simple as graph connectivity. We give an algorithm for dynamic graph connectivity in this setting with constant communication rounds and communication cost almost linear in terms of the batch size. Our techniques combine a new graph contraction technique, an independent random sample extractor from correlated samples, as well as distributed data structures supporting parallel updates and queries in batches. We also illustrate the power of dynamic algorithms in the MPC model by showing that the batched version of the adaptive connectivity problem is P-complete in the centralized setting, but sub-linear sized batches can be handled in a constant number of rounds. We believe our approach represents a practically-motivated workaround to the current difficulties in designing more efficient massively parallel static graph algorithms.

Laxman Dhulipala
Carnegie Mellon University

ldhulipa@cs.cmu.edu

David Durfee
LinkedIn
ddurfee@linkedin.com

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

Richard Peng, Saurabh Sawlani
Georgia Institute of Technology
richard.peng@gmail.com, sawlani@gatech.edu

Xiaorui Sun
University of Illinois at Chicago
xiaorui@uic.edu

CP22

Shortest Paths in a Hybrid Network Model

We introduce a communication model for *hybrid networks*, where nodes have access to two different communication modes: a *local* mode where (like in traditional networks) communication is only possible between specific pairs of nodes, and a *global* mode where (like in overlay networks) communication between any pair of nodes is possible. Typically, communication over short-range connections is cheaper and can be done at a much higher rate than communication via the overlay network. Therefore, we are focusing on the LOCAL model for the local connections. For the global communication we assume the *node-capacitated clique model*, where in each round every node can exchange $O(\log n)$-bit messages with $O(\log n)$ arbitrary nodes. We study the impact of hybrid communication on the complexity of distributed algorithms on the problem of computing shortest paths. For the all-pairs shortest paths problem, we show that an exact solution can be computed in time $\tilde{O}(n^{2/3})$. Approximate solutions can be computed in time $\tilde{\Theta}(n^{1/2})$ but not faster. The single-source shortest paths problem can be solved exactly in time $\tilde{O}(\sqrt{\mathsf{SPD}})$, where SPD denotes the shortest path diameter. A $(1+o(1))$-approximate solution can be computed in time $\tilde{O}(n^{1/3})$. Finally, for any constant $\varepsilon > 0$, we can compute an $O(1)$-approximate solution in time $\tilde{O}(n^{\varepsilon})$.

John Augustine
Indian Institute Of Technology Madras
j.e.augustine@gmail.com

Kristian Hinnenthal
University of Paderborn
krijan@mail.upb.de

Fabian Kuhn
University of Freiburg, Germany
kuhn@cs.uni-freiburg.de

Christian Scheideler
University of Paderborn
scheideler@upb.de

Philipp Schneider
University of Freiburg
philipp.schneider@cs.uni-freiburg.de

CP23

List Decoding of Direct Sum Codes

We consider families of codes obtained by "lifting" a base code $\mathcal{C}$ through operations such as k-XOR applied to "local views" of codewords of $\mathcal{C}$, according to a suitable k-uniform hypergraph. The k-XOR operation yields the direct sum encoding used in works of [Ta-Shma, STOC 2017] and [Dinur and Kaufman, FOCS 2017]. We give a general framework for list decoding such lifted codes, as long as the base code admits a unique decoding algorithm, and the hypergraph used for lifting satisfies certain expansion properties. We show that these properties are indeed satisfied by the collection of k-length walks on a sufficiently strong expanding graph, and by hypergraphs corresponding to high-dimensional expanders. Our framework relies on relaxations given by the Sum-of-Squares (SOS) SDP hierarchy for solving various constraint satisfaction problems (CSPs). We view the problem of recovering the closest codeword to a given word, as finding the optimal solution to an instance of a CSP. Constraints in the instance correspond to edges of the lifting hypergraph, and the solutions are restricted to lie in the base code $\mathcal{C}$. We extend the framework to list decoding, by requiring the SOS solution to minimize a convex proxy for negative entropy. We show that this ensures a covering property for the SOS solution, and the "condition and round" approach used in several SOS algorithms can then be used to recover the required list of codewords.

Vedat Levi Alev
University of Waterloo
vlalev@uwaterloo.ca

Fernando Granha Jeronimo, Dylan Quintana
University of Chicago
granha@uchicago.edu, dquintana@uchicago.edu

Shashank Srivastava
TTIC
shashanks@ttic.edu

Madhur Tulsiani
TTI Chicago
madhurt@ttic.edu

CP23

Relaxed Locally Correctable Codes with Nearly-Linear Block Length and Constant Query Complexity

Locally correctable codes (LCCs) are codes $C\colon \Sigma^k \to \Sigma^n$ which admit local algorithms that can correct any individual symbol of a corrupted *codeword* via a minuscule number of queries. One of the central problems in algorithmic coding theory is to construct $O(1)$-query LCC with minimal block length. Alas, state-of-the-art of such codes requires exponential block length to admit $O(1)$-query algorithms for local correction, despite much attention during the last two decades. This lack of progress prompted the study of *relaxed LCCs*, which allow the correction algorithm to *abort* (but not err) on small fraction of the locations. This relaxation turned out to allow constant-query correction algorithms for codes with polynomial block length. Specifically, prior work showed that there exist $O(1)$-query relaxed LCCs that achieve nearly-

quartic block length $n = k^{4+\alpha}$, for an arbitrarily small constant $\alpha > 0$. We construct an $O(1)$-query relaxed LCC with *nearly-linear* block length $n = k^{1+\alpha}$, for an arbitrarily small constant $\alpha > 0$. This significantly narrows the gap between the lower bound which states that there are no $O(1)$-query relaxed LCCs with block length $n = k^{1+o(1)}$. In particular, this resolves an open problem raised by Gur, Ramnarayan, and Rothblum (ITCS 2018).

Tom Gur
University of Warwick
tom.gur@warwick.ac.uk

CP23

On the Power of Relaxed Local Decoding Algorithms

A *locally decodable code* (LDC) $C\colon \{0,1\}^k \to \{0,1\}^n$ is an error correcting code that admits algorithms for recovering individual bits of the message by only querying a few bits of a noisy codeword. LDCs found a myriad of applications both in theory and in practice, ranging from probabilistically checkable proofs to distributed storage. However, despite nearly two decades of extensive study, the best known constructions of LDCs with $O(1)$-query decoding algorithms have super-polynomial blocklength. The notion of *relaxed LDCs* is a natural relaxation of LDCs, which aims to bypass the foregoing barrier by requiring local decoding of nearly all individual message bits, yet allowing decoding failure (but not error) on the rest. State of the art constructions of $O(1)$-query relaxed LDCs achieve blocklength $n = O\left(k^{1+\gamma}\right)$ for an arbitrarily small constant γ. Using algorithmic and combinatorial techniques, we prove an impossibility result, showing that codes with blocklength $n = k^{1+o(1)}$ cannot be relaxed decoded with $O(1)$-query algorithms. This resolves an open problem raised by Goldreich in 2004.

Tom Gur
University of Warwick
tom.gur@warwick.ac.uk

Oded Lachish
Birkbeck, University of London
oded@dcs.bbk.ac.uk

CP23

On the Performance of Reed-Muller Codes with Respect to Random Errors and Erasures

This work proves new results on the ability of binary Reed-Muller codes to decode from random errors and erasures. Specifically, we prove that RM codes with m variables and degree γm, for some explicit constant γ, achieve capacity for random erasures (i.e. for the binary erasure channel) and for random errors (for the binary symmetric channel). Earlier, it was known that RM codes achieve capacity for the binary symmetric channel for degrees $r = o(m)$. For the binary erasure channel it was known that RM codes achieve capacity for degree $o(m)$ or $r \in [m/2 \pm O(\sqrt{m})]$. Thus, our result provide a new range of parameters for which RM achieve capacity for these two well studied channels. In addition, our results imply that for every $\epsilon > 0$ (in fact we can get up to $\epsilon = \Omega\left(\frac{\sqrt{\log m}}{\sqrt{m}}\right)$) RM codes of degree $r < (1/2 - \epsilon)m$ can correct a fraction of $1 - o(1)$ random erasures with high probability. We also show that, information theoretically, such codes can handle a fraction of $\frac{1}{2} - o(1)$ random errors with high probability. These results show that RM codes of rates up to $1/\text{poly}(\log n)$ (where $n = 2^m$ is the block length) are is some sense as good as capacity achieving codes. We obtain these results by proving improved bounds on the weight distribution of Reed-Muller codes of high degrees.

Ori Sberlo, Amir Shpilka
Tel Aviv University
ori.sberlo@gmail.com, shpilka@tauex.tau.ac.il

CP23

On Decoding Cohen-Haeupler-Schulman Tree Codes

Tree codes, introduced by Schulman, are combinatorial structures essential to coding for interactive communication. An infinite family of tree codes with both rate and distance bounded by positive constants is called asymptotically good. Rate being constant is equivalent to the alphabet size being constant. Schulman proved that there are asymptotically good tree code families using the Lovasz local lemma, yet their explicit construction remains an outstanding open problem. In a major breakthrough, Cohen, Haeupler and Schulman constructed explicit tree code families with constant distance, but over an alphabet polylogarithmic in the length. Our main result is a randomized polynomial time decoding algorithm for these codes making novel use of the polynomial method. The number of errors corrected scales roughly as the block length to the three-fourths power, falling short of the constant fraction error correction guaranteed by the constant distance. We further present number theoretic variants of Cohen-Haeupler-Schulman codes, all correcting a constant fraction of errors with polylogarithmic alphabet size. Towards efficiently correcting close to a constant fraction of errors, we propose a speculative convex optimization approach inspired by compressed sensing.

Anand Kumar Narayanan
Sorbonne Universite
anand.narayanan@lip6.fr

Matthew Weidner
Carnegie Mellon University
maweidne@andrew.cmu.edu

CP24

Counting Independent Sets in Unbalanced Bipartite Graphs

Understanding the complexity of approximately counting the number of weighted or unweighted independent sets in a bipartite graph (#BIS) is a central open problem in the field of approximate counting. Here we consider a subclass of this problem and give an FPTAS for approximating the partition function of the hard-core model for bipartite graphs when there is sufficient imbalance in the degrees or fugacities between the sides (L, R) of the bipartition. This includes, among others, the biregular case when $\lambda = 1$ (approximating the number of independent sets of G) and $\Delta_R \geq 7\Delta_L \log(\Delta_L)$. Our approximation algorithm is based on truncating the cluster expansion of a polymer model partition function that expresses the hard-core partition function in terms of deviations from independent sets that are empty on one side of the bipartition. Further consequences of this method for unbalanced bipartite graphs include an efficient sampling algorithm for the hard-core model and zero-freeness results for the partition function with complex fugacities. By using connections between the cluster expansion and joint cumulants of certain random variables, we prove the hard-core model has exponential decay of correlations for all graphs and fugacities satisfying our conditions. This illustrates the applicability of statistical mechanics tools to algorithmic problems and refines our understanding of the connections between different methods of approximate counting.

Sarah Cannon
Claremont McKenna College
scannon@cmc.edu

Will Perkins
University of Illinois at Chicago
math@willperkins.org

CP24

Faster Sublinear Approximations of k-Cliques for Low Arboricity Graphs

Given query access to an undirected graph G, we consider the problem of computing a $(1 \pm \epsilon)$-approximation of the number of k-cliques in G. The standard query model for general graphs allows for degree queries, neighbor queries, and pair queries. Let n be the number of vertices, m be the number of edges, and n_k be the number of k-cliques. Previous work by Eden, Ron and Seshadhri (STOC 2018) gives an $O^*(\frac{n}{n_k^{1/k}} + \frac{m^{k/2}}{n_k})$-time algorithm for this problem (we use $O^*(\cdot)$ to suppress $(\log n, 1/\epsilon, k^k)$ dependencies). Moreover, this bound is nearly optimal when the expression is sublinear in the size of the graph. Our motivation is to circumvent this lower bound, by parameterizing the complexity in terms of *graph arboricity*. The arboricity of G is a measure for the graph density "everywhere'. There is a very rich family of graphs with bounded arboricity, including all minor-closed graph classes (such as planar graphs and graphs with bounded treewidth), bounded degree graphs, preferential attachment graphs and more. We design an algorithm for the class of graphs with arboricity at most α, whose running time is $O^*(\min\{\frac{n\alpha^{k-1}}{n_k}, \frac{n}{n_k^{1/k}} + \frac{m\alpha^{k-2}}{n_k}\})$. We also prove a nearly matching lower bound. For all graphs, the arboricity is $O(\sqrt{m})$, so this bound subsumes all previous results on sublinear clique approximation.

Talya Eden
Tel Aviv University
talyaa01@gmail.com

Dana Ron
Tel-Aviv University
danaron@tau.ac.il

C. Seshadhri
University of California, Santa Cruz
sesh@ucsc.edu

CP24

Ultimate Greedy Approximation of Independent Sets in Subcubic Graphs

We study the approximability of the maximum size independent set (MIS) problem in bounded degree graphs. This is one of the most classic and widely studied NP-hard and hard to approximate, optimization problems. We focus on the well known minimum degree greedy algorithm for MIS. The approximation ratios of this algorithm have been very widely studied, where it is augmented with an advice that tells the greedy which minimum degree vertex to choose if it is not unique. Our main contribution is a new mathematical theory for the design of such greedy algorithms with efficiently computable advice and for the analysis of their approximation ratios. With this new theory we obtain the ultimate approximation ratio of 5/4 for greedy on subcubic graphs, i.e., with maximum degree 3, which completely solves the open problem from the paper by Halldórsson and Yoshihara (1995). Our algorithm is the fastest currently known algorithm with this approximation ratio on such graphs. We also apply our techniques to prove bounds on the approximation ratios of greedy on graphs with any bounded maximum degree, and to the minimum vertex cover problem on subcubic graphs. We complement our positive, upper bound results with negative, lower bound results which prove that the problem of designing good advice for greedy is computationally or unconditionally hard, even to approximate, on various classes of graphs. These results significantly improve on such previously known hardness results.

Piotr J. Krysta
University of Liverpool
Department of Computer Science
pkrysta@liverpool.ac.uk

Mathieu Mari
Ecole Normale Superieure, Paris
mathieu.mari@ens.fr

Nan Zhi
University of Liverpool
existenzhi@gmail.com

CP24

Improved Inapproximability of Rainbow Coloring

A rainbow q-coloring of a k-uniform hypergraph is a q-coloring of the vertex set such that every hyperedge contains all q colors. We prove that given a rainbow $(k - 2\lfloor\sqrt{k}\rfloor)$-colorable k-uniform hypergraph, it is NP-hard to find a normal 2-coloring. Previously, this was only known for rainbow $\lfloor k/2\rfloor$-colorable hypergraphs (Guruswami and Lee, SODA 2015). We also study a generalization which we call rainbow (q, p)-coloring, defined as a coloring using q colors such that every hyperedge contains at least p colors. We prove that given a rainbow $(k - \lfloor\sqrt{kc}\rfloor, k - \lfloor 3\sqrt{kc}\rfloor)$-colorable k uniform hypergraph, it is NP-hard to find a normal c-coloring for any $c = o(k)$. The proof of our second result relies on two combinatorial theorems. One of the theorems was proved by Sarkaria (J. Comb. Theory, Ser. B 1990) using topological methods and the other theorem we prove using a generalized Borsuk-Ulam theorem.

Aditya Potukuchi
Rutgers University
apotu.57@gmail.com

Per Austrin
KTH Royal Institute of Technology
austrin@kth.se

Amey Bhangale
Weizmann Institute

ameyrbh@gmail.com

CP24

Improved Hardness for H-Colourings of G-Colourable Graphs

We present new results on approximate colourings of graphs and, more generally, approximate H-colourings and promise constraint satisfaction problems. First, we show NP-hardness of colouring k-colourable graphs with $\binom{k}{k/2}-1$ colours for every $k \geq 4$. This improves the result of Buln, Krokhin, and Opral [STOC'19], who gave NP-hardness of colouring k-colourable graphs with $2k-1$ colours for $k \geq 3$, and the result of Huang [APPROX-RANDOM'13], who gave NP-hardness of colouring k-colourable graphs with $2^{k^{1/3}}$ colours for sufficiently large k. Thus we improve from known linear/sub-exponential gaps to exponential gaps. Second, we show that the topology of the box complex of H alone determines whether H-colouring of G-colourable graphs is NP-hard for all (non-bipartite, H-colourable) G. This formalises the topological intuition behind the result of Krokhin and Opral [FOCS'19] that 3-colouring of G-colourable graphs is NP-hard for all (3-colourable, non-bipartite) G. We use this technique to establish NP-hardness of H-colouring of G-colourable graphs for H that include but go beyond K_3, including square-free graphs and circular cliques (leaving K_4 and larger cliques open). Underlying all of our proofs is a very general observation that adjoint functors give reductions between promise constraint satisfaction problems.

Marcin Wrochna
University of Oxford, UK
marcin.wrochna@cs.ox.ac.uk

Stanislav Zivny
University of Oxford
standa.zivny@cs.ox.ac.uk

CP25

Distributionally Robust Max Flows

We study a distributionally robust max flow problem under the marginal distribution model, where the vector of arc capacities is random, with the marginals to the joint multivariate distribution being known, but the correlation being unknown. The goal is to compute the expected value of the max flow under the worst-case joint distribution of arc capacities. We provide a simple combinatorial proof that shows that for the case of finite-supported marginal distributions, this worst-case expectation can be efficiently computed, and moreover, the worst-case joint distribution can be explicitly constructed, despite being non-trivial in the sense that it is not a combination of monotonic or anti-monotonic couplings. Our technique is to use a related min-cost flow problem to generate a distribution over cuts in the graph, which in turn induces the worst-case joint distribution. It also provides an alternative interpretation of the problem as a zero-sum game between a capacity player and a cut player.

Louis L. Chen
Naval Postgraduate School
louischenusa@gmail.com

Will Ma
Columbia University
wm2428@columbia.edu

James Orlin, David Simchi-Levi
Massachusetts Institute of Technology
jorlin@mit.edu, dslevi@mit.edu

CP25

On a Decentralized $(\Delta{+}1)$-Graph Coloring Algorithm

We consider a decentralized graph coloring model where each vertex only knows its own color and whether some neighbor has the same color as it. The networking community has studied this model extensively due to its applications to channel selection, rate adaptation, etc. Here, we analyze variants of a simple algorithm of Bhartia et al. [Proc., ACM MOBIHOC, 2016]. In particular, we introduce a variant which requires only $O(n \log \Delta)$ expected recolorings that generalizes the coupon collector problem. Finally, we show that the $O(n\Delta)$ bound Bhartia et al. achieve for their algorithm still holds and is tight in adversarial scenarios.

Paul B. De Supinski
Facebook
pdesupinski@yahoo.com

Deeparnab Chakrabarty
Dartmouth College
deeparnab@dartmouth.edu

CP25

Simple Label-Correcting Algorithms for Partially-Dynamic Approximate Shortest Paths in Directed Graphs

Classical single-source shortest paths algorithms work by maintaining distance estimates d and performing so-called edge relaxations. We call an edge uv of weight $w(uv)$ relaxed if $d(v) \leq d(u) + w(uv)$, and tense otherwise. To relax a tense edge uv means to set $d(v)$ to $d(u) + w(uv)$. It is known that starting from $d(s) = 0$, and $d(v) = \infty$ for all $v \neq s$, and performing relaxations until there are no tense edges leads to d being equal to the distances from the source s. This overall idea can be extended to a simple incremental shortest paths algorithm. We consider an operation which can be seen as a dual of a relaxation and study an approximate version of both operations. We show that by repeating the respective operation until convergence one obtains simple incremental and decremental deterministic algorithms for $(1+\epsilon)$-approximate shortest paths in directed graphs. Our algorithm maintaining all-pairs approximate shortest paths has $O(n^3 \log(n) \log(nW)/\epsilon)$ total update time if the edge weights come from the set $[1, W]$. This is two log-factors faster than the known solution obtained by combining King's decremental transitive closure algorithm [FOCS'99] and Bernstein's h-SSSP algorithm [SICOMP'16] for $h = 2$. In addition, we give an algorithm for approximating single-source shortest paths of hop-length at most h in $O(mh \log(nW)/\epsilon)$ total time. The obtained algorithm is simpler and more efficient than the h-SSSP algorithm.

Adam Karczmarz
Institute of Informatics, University of Warsaw
a.karczmarz@mimuw.edu.pl

Jakub Lacki
Google Research

jlacki@google.com

CP25

Distributed Backup Placement in One Round and its Applications to Maximum Matching Approximation and Self-Stabilization

In the distributed backup-placement problem each node of a network has to select one neighbor, such that the maximum number of nodes that make the same selection is minimized. This is a natural relaxation of the perfect matching problem, in which each node is selected just by one neighbor. Previous (approximate) solutions for backup placement are non-trivial, even for simple graph topologies, such as dense graphs. In this paper we devise an algorithm for dense graph topologies. Our algorithm requires just one round, and the approximation ratio is improved by a multiplicative factor of at least 2. Our new algorithm has several interesting implications. In particular, it gives rise to a $(2+\epsilon)$-approximation to maximum matching within $O(\log^* n)$ rounds in dense networks. The resulting algorithm is very simple as well, in sharp contrast to previous algorithms that compute such a solution within this running time. Moreover, these algorithms are applicable to a narrower graph family than our algorithm. For the same graph family, the best previously-known result has $O(\log \Delta + \log^* n)$ running time. Another interesting implication is the possibility to execute our backup placement algorithm as-is in the self-stabilizing setting. This makes it possible to simplify and improve other algorithms for the self-stabilizing setting, by employing helpful properties of backup placement.

Gal Oren
Ben-Gurion University of the Negev
Nuclear Research Center - Negev
orenw@post.bgu.ac.il

Leonid Barenboim
Open University of Israel
leonidb@openu.ac.il

CP25

Nearly Linear Time Approximations for Mixed Packing and Covering Problems Without Data Structures Or Randomization

We consider nearly linear time approximations for explicit mixed packing and covering LPs, a class of linear programs with many applications. Young [2014] gave a $O(N \log(m)/\epsilon^2)$ deterministic algorithm based on the multiplicative weight update (MWU) framework, where N is the total number of nonzeroes, m is the number of constraints, and ϵ is the relative error parameter. The key technical ingredient was a data structure for managing the "multiplicative weight updates' more efficiently, which is nontrivial. One can also obtain the same running time by randomly sampling the multiplicative weight updates [Chekuri and Quanrud, 2018], but the probabilistic analysis is technical. We give a different deterministic algorithm that runs in $(N \log(m)(\log\log(n)+\log(1/\epsilon))/\epsilon^2)$ time. The new algorithm is also based on the MWU framework, and simply replaces the so-called "width-independent step size' with a maximal step size found by binary search; a.k.a. a line search. It does not require data structures or randomization to manage the weights; one just recomputes the weights as needed.

Kent Quanrud
University of Illinois at Urbana-Champaign
krq@purdue.edu

CP26

Chasing Convex Bodies with Linear Competitive Ratio

We study the problem of chasing convex bodies online: given a sequence of convex bodies $K_t \subseteq \mathbb{R}^d$ the algorithm must respond with points $x_t \in K_t$ in an online fashion (i.e., x_t is chosen before K_{t+1} is revealed). The objective is to minimize the total distance between successive points in this sequence. Recently, Bubeck et al. (STOC 2019) gave a $2^{O(d)}$-competitive algorithm for this problem. We give an algorithm that is $O(\min(d, \sqrt{d \log T}))$-competitive for any sequence of length T.

C.J. Argue, Anupam Gupta
Carnegie Mellon University
cargue@andrew.cmu.edu, anupamg@cs.cmu.edu

Guru Guruganesh
Google Research
prashanth818@gmail.com

Ziye Tang
Carnegie Mellon University
ziyet@andrew.cmu.edu

CP26

Online Probabilistic Metric Embedding: A General Framework for Bypassing Inherent Bounds

Probabilistic metric embedding into trees is a powerful tool for designing online algorithms. The standard approach is to embed the entire metric into a tree metric and then solve the problem on the latter. The overhead in the competitive ratio depends on the expected distortion of the embedding, which is logarithmic in the size of the underlying metric. For many online applications it is natural to ask if it is possible to construct such embeddings in an online fashion such that the distortion would be a polylogarithmic function of k, the number of requested points. We answer this question negatively, showing a lower bound of $\tilde{\Omega}(\log k \log r)$, where r is the aspect ratio of the set of requested points, showing that a modification of the probabilistic embedding into trees of Bartal (FOCS'96), which has expected distortion of $O(\log k \log r)$, is nearly-tight. Surprisingly, we develop a general framework for bypassing this limitation. We show that for a large class of online problems this online embedding can still be used to devise an algorithm with $O(\min\{\log k \log(k\lambda), \log^3 k\})$ overhead in the competitive ratio, where k is the current number of requested points, and λ is a measure of subadditivity of the cost function, which is at most r, the current number of requests. This implies the first algorithms with competitive ratio $polylog(k)$ for online subadditive network design and $polylog(k, r)$ for online group Steiner forest.

Yair Bartal
Hebrew University
yair@cs.huji.ac.il

Nova Fandina
Hebrew University of Jerusalem
fandina@cs.huji.ac.il

William Umboh

The University of Sydney
Australia
william.umboh@sydney.edu.au

CP26

Achieving Optimal Backlog in the Vanilla Multi-Processor Cup Game

In each step of the p-processor cup game on n cups, a filler distributes up to p units of water among the cups, subject only to the constraint that no cup receives more than 1 unit of water; an emptier then removes up to 1 unit of water from each of p cups. Designing strategies for the emptier that minimize backlog (i.e., the height of the fullest cup) is important for applications in processor scheduling, buffer management in networks, quality of service guarantees, and deamortization. We prove that the greedy algorithm (i.e., the empty-from-fullest-cups algorithm) achieves backlog $O(\log n)$ for any $p \geq 1$. This resolves a long-standing open problem for $p > 1$, and is asymptotically optimal as long as $n \geq 2p$. If the filler is an oblivious adversary, then we prove that there is a randomized emptying algorithm that achieve backlog $O(\log p + \log \log n)$ with probability $1 - 2^{-polylog(n)}$ for $2^{polylog(n)}$ steps. This is known to be asymptotically optimal when n is sufficiently large relative to p. Previously, the only known bound on backlog for $p > 1$, and the only known randomized guarantees for any p (including when $p = 1$), required the use of resource augmentation, meaning that the filler can only distribute at most $p(1 - \epsilon)$ units of water in each step, and that the emptier is then permitted to remove $1 + \delta$ units of water from each of p cups, for some $\epsilon, \delta > 0$.

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

CP26

The Online Submodular Cover Problem

Motivated by problems in network monitoring and resource allocation, we consider the submodular cover problem in an online setting. At each time t, a nonnegative monotone submodular function g_t is given to us. We define $f^{(t)} = \sum_{s \leq t} g_s$ as the sum of all functions seen so far. We want to maintain a min-cost submodular cover of these submodular functions $f^{(1)}, f^{(2)}, \ldots f^{(T)}$ in an online fashion; i.e., we cannot revoke previous choices. Formally, at each time t we produce a set $S_t \subseteq N$ such that $f^{(t)}(S_t) = f^{(t)}(N)$ and $S_{t-1} \subseteq S_t$. We give polylogarithmic competitive algorithms for this online submodular cover problem. The competitive ratio on an input sequence of length T is $O(\ln n \ln(T \cdot f_{\max}/f_{\min}))$, where $f_{\max}$ and $f_{\min}$ are the largest and smallest marginals for functions $f^{(t)}$, and $|N| = n$. For the special case of online set cover, our competitive ratio matches that of (Alon et al. 03), which are best possible for polynomial-time online algorithms unless $NP \subseteq BPP$ (Korman 04). The technical challenge is to (approximately) solve Wolsey's exponential-sized linear programming relaxation for submodular cover, and to round it, both in the online setting. Moreover, to get our competitiveness bounds, we define a (seemingly new) generalization of mutual information to general submodular functions, which we call mutual coverage.

Roie Levin, Anupam Gupta
Carnegie Mellon University
roiel@andrew.cmu.edu, anupamg@cs.cmu.edu

CP26

Chasing Convex Bodies Optimally

In the chasing convex bodies problem, an online player receives a request sequence of N convex sets $K_1, \ldots, K_N$ contained in a normed space $\mathbb{R}^d$. The player starts at $x_0 \in \mathbb{R}^d$, and after observing each K_n picks a new point $x_n \in K_n$. At each step the player pays a movement cost of $||x_n - x_{n-1}||$. The player aims to maintain a constant competitive ratio against the minimum cost possible in hindsight, i.e. knowing all requests in advance. The existence of a finite competitive ratio for convex body chasing was first conjectured in 1991 by Friedman and Linial. This conjecture was recently resolved in a paper which proved an exponential $2^{O(d)}$ upper bound on the competitive ratio. In this paper, we drastically improve the exponential upper bound. We give an algorithm achieving competitive ratio d for arbitrary normed spaces, which is exactly tight for ℓ^∞. In Euclidean space, our algorithm achieves nearly optimal competitive ratio $O(\sqrt{d \log N})$, compared to a lower bound of $\sqrt{d}$. Our approach extends our recent work which chases nested convex bodies using the classical Steiner point of a convex body. We define the functional Steiner point of a convex function and apply it to the work function to obtain our algorithm.

Mark Sellke
Stanford University
msellke@gmail.com

CP26

Chasing Nested Convex Bodies Nearly Optimally

The convex body chasing problem, introduced by Friedman and Linial, is a competitive analysis problem on any normed vector space. In convex body chasing, for each timestep $t \in \mathbb{N}$, a convex body $K_t \subseteq \mathbb{R}^d$ is given as a request, and the player picks a point $x_t \in K_t$. The player aims to ensure that the total distance $\sum_{t=0}^{T-1} ||x_t - x_{t+1}||$ is within a bounded ratio of the smallest possible offline solution. In this work, we consider the nested version of the problem, in which the sequence (K_t) must be decreasing. For Euclidean spaces, we consider a memoryless algorithm which moves to the so-called Steiner point, and show that in a certain sense it is exactly optimal among memoryless algorithms. For general finite dimensional normed spaces, we combine the Steiner point and our recent algorithm to obtain an algorithm which is nearly optimal for all ℓ^p_d spaces with $p \geq 1$, closing a polynomial gap.

Mark Sellke
Stanford University
msellke@gmail.com

Yin Tat Lee
University of Washington
yintat@uw.edu

Sébastien Bubueck
Microsoft Research
sebubeck@microsoft.com

Yuanzhi Li
Carnegie Mellon University
yuanzhil@stanford.edu

Bo'az Klartag
Weizmann Institute
klartag@gmail.com

CP27

Adaptive Shivers Sort: An Alternative Sorting Algorithm

We present a new sorting algorithm, called adaptive ShiversSort, that exploits the existence of monotonic runs for sorting efficiently partially sorted data. This algorithm is a variant of the well-known algorithm TimSort, which is the sorting algorithm used in standard libraries of programming languages such as Python or Java (for non-primitive types). More precisely, adaptive ShiversSort is a so-called k-aware merge-sort algorithm, a class that was introduced by Buss and Knop that captures "TimSort-like' algorithms. In this article, we prove that, although adaptive ShiversSort is simple to implement and differs only slightly from TimSort, its computational cost, in number of comparisons performed, is optimal within the class of *natural* merge-sort algorithms, up to a small additive linear term: this makes adaptive ShiversSort the first k-aware algorithm to benefit from this property, which is also a 33% improvement over TimSort's worst-case. This suggests that adaptive ShiversSort could be a strong contender for being used instead of TimSort. Then, we investigate the optimality of k-aware algorithms: we give lower and upper bounds on the best approximation ratios of such algorithms, compared to optimal stable natural merge-sort algorithms. In particular, we design generalisations of adaptive ShiversSort whose computational costs are optimal up to arbitrarily small multiplicative factors.

Vincent Jugé
Université Paris-Est Marne-la-Vallée
vincent.juge@univ-mlv.fr

CP27

Hyperbolic Intersection Graphs and (quasi)-Polynomial Time

We study intersection graphs of unit-ball-like objects in hyperbolic space. In the hyperbolic plane, we introduce a new technique to bound the treewidth of these graphs. The treewidth bounds yield quasi-polynomial ($n^{O(logn)}$) algorithms for Independent Set and several other problems, while in the case of Hamiltonian Cycle and 3-Coloring we even get polynomial time algorithms. Then we will show that the quasi-polynomial algorithm for Independent Set is optimal up to constant factors in the exponent under ETH. I will also talk about a separator for intersection graphs in higher dimensional hyperbolic space, which leads to further ETH-tight algorithms. The running times of these algorithms in d-dimensional hyperbolic space (for $d \geq 3$) are the same as the running times of similar algorithms in $d-1$-dimensional Euclidean space.

Sándor Kisfaludi-Bak
Eindhoven University of Technology
sandor.kisfaludi-bak@mpi-inf.mpg.de

CP27

Fine-Grained Complexity of Graph Homomorphism Problem for Bounded-Treewidth Graphs

Abstract not available.

Karolina Okrasa, Pawel Rzazewski
Warsaw University of Technology, Poland
k.okrasa@mini.pw.edu.pl, p.rzazewski@mini.pw.edu.pl

CP27

Improved Algorithms for Edit Distance and LCS: Beyond Worst Case

Edit distance and longest common subsequence are among the most fundamental problems in combinatorial optimization. Recent developments have proven strong lower bounds against subquadratic time solutions for both problems. Moreover, the best approximation factors for subquadratic time solutions have been limited to 3 for edit distance and super constant for longest common subsequence. Improved approximation algorithms for these problems[1] are some of the biggest open questions in combinatorial optimization. In this work, we present improved algorithms for both edit distance and longest common subsequence. The running times are truly subquadratic, though we obtain $1 + o(1)$ approximate solutions for both problems if the input satisfies a mild condition. In this setting, first, an adversary chooses one of the input strings. Next, this string is perturbed by a random procedure and then the adversary chooses the second string **after observing the perturbed one**.

Mahdi Boroujeni, Masoud Seddighin
Sharif University of Technology
safarnejad@gmail.com, seddighin.masood@gmail.com

Saeed Seddighin
University of Maryland, College Park
saeedreza.seddighin@gmail.com

CP27

Reducing Approximate Longest Common Subsequence to Approximate Edit Distance

Given a pair of n-character strings, the problems of computing their Longest Common Subsequence and Edit Distance have been extensively studied for decades. For exact algorithms, LCS and Edit Distance (with character insertions and deletions) are equivalent; the state of the art running time is (almost) quadratic in n, and this is tight under plausible fine-grained complexity assumptions. But for approximation algorithms the picture is different: there is a long line of works with improved approximation factors for Edit Distance, but for LCS (with binary strings) only a trivial 1/2-approximation was known. In this work we give a reduction from approximate LCS to approximate Edit Distance, yielding the first efficient $(1/2+\epsilon)$-approximation algorithm for LCS for some constant $\epsilon > 0$.

Aviad Rubinstein
Harvard University
aviad@cs.stanford.edu

Zhao Song
UT-Austin

[1] $1+\epsilon$ approximation for ED or subpolynomial approximation for LCS in truly subquadratic time.

zhaos@utexas.edu

CP28

The Communication Complexity of Set Intersection and Multiple Equality Testing

In this work, We explore fundamental problems in randomized communication complexity such as computing Set Intersection on sets of size k and Equality Testing between vectors of length k. It has been established that one can achieve optimal communication complexity $O(k)$ with a randomized protocol that takes $\log^* k$ rounds, and this is one point in the optimal round-communication tradeoff curve. Besides the round-communication tradeoff, there is a third parameter of interest, namely the error probability. In this work, we show that:

- Any protocol for solving Multiple Equality Testing in r rounds with failure probability $= 2^{-E}$ has communication volume $\Omega(Ek^{1/r})$.
- There exists a protocol for solving Multiple Equality Testing in $r + \log^*(k/E)$ rounds with $O(k + rEk^{1/r})$ communication, thereby essentially matching our lower bound and previous lower bounds.
- Lower bounds on Equality Testing extend to Set Intersection, for every r, k, and (which is trivial); in the reverse direction, upper bounds on Equality Testing for r, k, imply similar upper bounds on Set Intersection with parameters $r + 1, k$, and .

Dawei Huang, Seth Pettie
University of Michigan
dwhuang9@gmail.com, pettie@seth.net

Zhijun Zhang
Institute for Interdisciplinary Information Sciences
Tsinghua University
zhijun-z16@mails.tsinghua.edu.cn

Yixiang Zhang
IITS Tsinghua University
zhangyix16@mails.tsinghua.edu.cn

CP28

New (α, β) Spanners and Hopsets

An $f(d)$-spanner of an unweighted n-vertex graph $G = (V, E)$ is a subgraph H satisfying that $_H(u, v)$ is at most $f(_G(u, v))$ for every $u, v \in V$. A simple girth argument implies that any $f(d)$-spanner with $O(n^{1+1/k})$ edges must satisfy that $f(d)/d = \Omega(\lceil k/d \rceil)$. A matching upper bound (even up to constants) for super-constant values of d is currently known only for $d = \Omega((\log k)^{\log k})$ as given by the well known $(1 + \epsilon, \beta)$ spanners of Elkin and Peleg, and its recent improvements by [Elkin-Neiman, SODA'17], and [Abboud-Bodwin-Pettie, SODA'18]. We present new spanner constructions that achieve a nearly optimal stretch of $O(\lceil k/d \rceil)$ for any distance value $d \in [1, k^{1-o(1)}]$, and $d \geq k^{1+o(1)}$. We also consider the related graph concept of hopsets introduced by [Cohen, J. ACM '00]. Informally, an hopset H is a weighted edge set that, when added to the graph G, allows one to get a path from each node u to a node v with at most β hops (i.e., edges) and length at most $\alpha \cdot_G (u, v)$. Hopsets constructions are known currently only for a narrow set of parameters. We present a new family of (α, β) hopsets with $\widetilde{O}(k \cdot n^{1+1/k})$ edges and $\alpha \cdot \beta = O(k)$.

Merav Parter, Uri Ben-Levy
Weizmann Institute
merav.parter@weizmann.ac.il,
uri.benlevy@weizmann.ac.il

CP28

Composable Core-Sets for Determinant Maximization Problems Via Spectral Spanners

We study a generalization of classical combinatorial graph spanners to the spectral setting. Given a set of vectors $V \subseteq \mathbb{R}^d$, we say a set $U \subseteq V$ is an α-spectral k-spanner, for $k \leq d$, if for all $v \in V$ there is a probability distribution μ_v supported on U such that

$$vv\overset{\preceq}{_k}\alpha \cdot \mathbb{E}_{u\sim\mu_v} uu'$$

where for two matrices $A, B \in \mathbb{R}^{d\times d}$ we write $A \preceq_k B$ iff the sum of the bottom $d - k + 1$ eigenvalues of $B - A$ is nonnegative. We show that any set V has an $\tilde{O}(k)$-spectral spanner of size $\tilde{O}(k)$ and this bound is almost optimal in the worst case. We use spectral spanners to study composable core-sets for spectral problems. We show that for many objective functions one can use a spectral spanner, independent of the underlying function, as a core-set and obtain almost optimal composable core-sets. For example, for the k-determinant maximization problem, we obtain an $\tilde{O}(k)^k$-composable core-set, and we show that this is almost optimal in the worst case. Our algorithm is a spectral analogue of the classical greedy algorithm for finding (combinatorial) spanners in graphs. We expect that our spanners find many other applications in distributed or parallel models of computation.

Piotr Indyk
Massachusetts Institute of Technology
indyk@mit.edu

Sepideh Mahabadi
Toyota Technological Institute at Chicago(TTIC)
sepideh.mahabadi@gmail.com

Shayan Oveis Gharan
University of Washington
shayan@cs.washington.edu

Alireza Rezaei
UNIVERSITY OF WASHINGTON
arezaei@cs.washington.edu

CP28

Tight Bounds for the Subspace Sketch Problem with Applications

In the subspace sketch problem one is given an $n \times d$ matrix A with $O(\log(nd))$ bit entries, and would like to compress it in an arbitrary way to build a small space data structure Q_p, so that for any given vector x, with probability at least 2/3, one has $Q_p(x) = (1 \pm \varepsilon)\|Ax\|_p$, where $p \geq 0$ and the randomness is over the construction of Q_p. The central question is: How many bits are necessary to store Q_p? We show if $p \geq 0$ is *not a positive even integer* and $d = \Omega(\log(1/\varepsilon))$, then $\widetilde{\Omega}(\varepsilon^{-2} \cdot d)$ bits are necessary. On the other hand, if p is a positive even integer, then there is an upper bound of $O(d^p \log(nd))$ bits independent of ε. Our results are optimal up to logarithmic factors, and show in particular that one cannot compress A to $O(d)$ "directions' $v_1, \ldots, v_{O(d)}$, such that for any x, $\|Ax\|_1$ can be well-approximated from $\langle v_1, x\rangle, \ldots, \langle v_{O(d)}, x\rangle$. Our lower bound

rules out arbitrary functions of these inner products (and in fact arbitrary data structures built from A), and thus rules out the possibility of a singular value decomposition for ℓ_1 in a very strong sense. Indeed, as $\varepsilon \to 0$, for $p = 1$ the space complexity becomes arbitrarily large, while for $p = 2$ it is at most $O(d^2 \log(nd))$.

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

Ruosong Wang, David Woodruff
Carnegie Mellon University
ruosongw@andrew.cmu.edu, dwoodruf@cs.cmu.edu

CP28

The Communication Complexity of Optimization

We consider the communication complexity of a number of distributed optimization problems. We start with the problem of solving a linear system. Suppose there is a coordinator together with s servers $P_1, \ldots, P_s$, the i-th of which holds a subset $A^{(i)}x = b^{(i)}$ of n_i constraints of a linear system in d variables, and the coordinator would like to output a vector x for which $A^{(i)}x = b^{(i)}$ for $i = 1, \ldots, s$. We assume each coefficient of each constraint is specified using L bits. We resolve the randomized and deterministic communication complexity, showing it is $\tilde{\Theta}(d^2L + sd)$ and $\tilde{\Theta}(sd^2L)$, respectively. When there is no solution to the linear system, a natural alternative is to find the solution minimizing the ℓ_p loss. While this problem has been studied, we give improved bounds for every value of $p \geq 1$. One takeaway message is that sampling and sketching techniques are neither optimal in the dependence on d nor on the dependence on the approximation ε, thus motivating new techniques from optimization. Towards this end, we consider the communication complexity of optimization tasks which generalize linear systems. For linear programming, we first resolve the communication complexity when d is constant, showing it is $\tilde{\Theta}(sL)$. For general d, we show an $\tilde{O}(sd^3L)$ upper bound and an $\tilde{\Omega}(d^2L+sd)$ lower bound.

Santosh Vempala
Georgia Tech
vempala@cc.gatech.edu

Ruosong Wang, David Woodruff
Carnegie Mellon University
ruosongw@andrew.cmu.edu, dwoodruf@cs.cmu.edu

CP29

Faster Update Time for Turnstile Streaming Algorithms

In this paper, we present a new algorithm for maintaining linear sketches in turnstile streams with faster update time. As an application, we show that $\log n$ `Count` sketches or `CountMin` sketches with a constant number of columns (i.e., buckets) can be implicitly maintained in *worst-case* $O(\log^{0.582} n)$ update time using $O(\log n)$ words of space, on a standard word RAM with word-size $w = \Theta(\log n)$. The exponent $0.582 \approx 2\omega/3 - 1$. Due to the numerous applications of linear sketches, our algorithm improves the update time for many streaming problems in turnstile streams, in the high success probability setting, without using more space, including ℓ_2 norm estimation, ℓ_2 heavy hitters, point query with ℓ_1 or ℓ_2 error, etc. Our algorithm generalizes, with the same update time and space, to maintaining $\log n$ linear sketches, where each sketch

1. partitions the coordinates into $k < \log^{o(1)} n$ buckets using a c-wise independent hash function for constant c,
2. maintains the sum of coordinates for each bucket.

Moreover, if arbitrary word operations are allowed, the update time can be further improved to $O(\log^{0.187} n)$, where $0.187 \approx \omega/2 - 1$. Our update algorithm is adaptive, and it circumvents the non-adaptive cell-probe lower bounds for turnstile streaming algorithms by Larsen, Nelson and Nguyên (STOC'15).

Josh Alman
MIT
jalman@seas.harvard.edu

Huacheng Yu
Princeton University
yuhch123@gmail.com

CP29

Approximate Maximum Matching in Random Streams

In this paper, we study the problem of finding a maximum matching in the semi-streaming model when edges arrive in a random order. A recent inspiring work by Assadi et al. shows that there exists a streaming algorithm with the approximation ratio of $\frac{2}{3}$ that uses $\tilde{O}(n^{1.5})$ memory. However, the memory of their algorithm is much larger than the memory constraint of the semi-streaming algorithms. In this work, we further investigate this problem in the semi-streaming model, and we present simple and clean algorithms for approximating maximum matching in the semi-streaming model. Our main results are as follows.

- We show that there exists a single-pass deterministic semi-streaming algorithm that finds a $\frac{3}{5}(= 0.6)$ approximation of the maximum matching in bipartite graphs using $\tilde{O}(n)$ memory. This result significantly outperforms the state-of-the-art result of Konrad that finds a 0.539 approximation of the maximum matching using $\tilde{O}(n)$ memory.
- By giving a black-box reduction from finding a matching in general graphs to finding a matching in bipartite graphs, we show there exists a single-pass deterministic semi-streaming algorithm that finds a $\frac{6}{11}(\approx 0.545)$ approximation of the maximum matching in general graphs, improving upon the state-of-art result 0.506 approximation by Gamlath et al.

Alireza Farhadi
University of Maryland
farhadi94@gmail.com

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

Tung Mai, Anup Rao, Ryan Rossi
Adobe Research
tumai@adobe.com, anuprao@adobe.com, rrossi@adobe.com

CP29

Vertex Ordering Problems in Directed Graph

Streams

We consider *directed* graph algorithms in a streaming setting, focusing on problems concerning orderings of the vertices. This includes such fundamental problems as topological sorting and acyclicity testing. We also study the related problems of finding a minimum feedback arc set (edges whose removal yields an acyclic graph), and finding a sink vertex. We are interested in both adversarially-ordered and randomly-ordered streams. For arbitrary input graphs with edges ordered adversarially, we show that most of these problems have high space complexity, precluding sublinear-space solutions. Some lower bounds also apply when the stream is randomly ordered: e.g., in our most technical result we show that testing acyclicity in the p-pass random-order model requires roughly $n^{1+1/p}$ space. For other problems, random ordering can make a dramatic difference: e.g., it is possible to find a sink in an acyclic tournament in the one-pass random-order model using polylog(n) space whereas under adversarial ordering roughly $n^{1/p}$ space is necessary and sufficient given $\Theta(p)$ passes. We also design sublinear algorithms for the feedback arc set problem in tournament graphs; for random graphs; and for randomly ordered streams. In some cases, we give lower bounds establishing that our algorithms are essentially space-optimal. Together, our results complement the much maturer body of work on algorithms for *undirected* graph streams.

Amit Chakrabarti, Prantar Ghosh
Dartmouth College
ac@cs.dartmouth.edu, prantar.ghosh.gr@dartmouth.edu

Andrew McGregor
University of Massachusetts, Amherst
mcgregor@cs.umass.edu

Sofya Vorotnikova
Dartmouth College
sofya.vorotnikova@dartmouth.edu

CP29

Fast and Space Efficient Spectral Sparsification in Dynamic Streams

In this paper, we resolve the complexity of the problem of spectral graph sparcification in dynamic streams up to polylogarithmic factors. Using a *linear sketch* we design a streaming algorithm that uses $\tilde{O}(n)$ space, and with high probability, recovers a spectral sparsifier from the sketch in $\tilde{O}(n)$ time. Prior results either achieved near optimal $\tilde{O}(n)$ space, but $\Omega(n^2)$ recovery time [Kapralov et al. '14], or ran in $o(n^2)$ time, but used polynomially suboptimal space [Ahn et al '13]. Our main technical contribution is a novel method for recovering graph edges with high effective resistance from a linear sketch. We show how to do so in nearly linear time by 'bucketing' vertices of the input graph into clusters using a coarse approximation to the graph's effective resistance metric. A second main contribution is a new pseudorandom generator (PRG) for linear sketching algorithms. Constructed from a *locally computable randomness extractor*, our PRG stretches a seed of $\tilde{O}(n)$ random bits polynomially in length with just $\log^{O(1)} n$ runtime cost per evaluation. This improves on Nisan's commonly used PRG, which in our setting would require $\tilde{O}(n)$ time per evaluation. Our faster PRG is essential to simultaneously achieving near optimal space and time complexity.

Jakab Tardos, Michael Kapralov, Aida Mousavifar
EPFL
jakab.tardos@epfl.ch, michael.kapralov@epfl.ch, aidasadat.mousavifar@epfl.ch

Cameron Musco
University of Massachusetts Amherst
cmusco@cs.umass.edu

Christopher Musco
New York University
cmusco@nyu.edu

Navid Nouri
EPFL
navid.nouri@epfl.ch

Aaron Sidford
Stanford University
sidford@standford.edu

CP29

Space Efficient Approximation to Maximum Matching Size from Uniform Edge Samples

Given a source of iid samples of edges of an input graph G with n vertices and m edges, how many samples does one need to compute a constant factor approximation to the maximum matching size in G? Moreover, is it possible to obtain such an estimate in a small amount of space? We show that this problem cannot be solved using a nontrivially sublinear (in m) number of samples: $m^{1-o(1)}$ samples are needed. On the other hand, $O(\log^2 n)$ bits of space suffice to compute an estimate. Our main technical tool is a new peeling type algorithm for matching and its local simulation. We show that a delicate balance between exploration depth and sampling rate allows our simulation to not lose precision over a logarithmic number of levels of recursion and achieve a constant factor approximation. Our algorithm also yields a constant factor approximate local computation algorithm (LCA) for matching with $O(d \log n)$ exploration starting from any vertex. Previous approaches were based on local simulations of randomized greedy, which take $O(d)$ time *in expectation over the starting vertex or edge* (Yoshida et al'09, Onak et al'12), and could not achieve a better than d^2 runtime. Interestingly, we also show that unlike our algorithm, the local simulation of randomized greedy that is the basis of the most efficient prior results does take $\Omega(d^2) \gg O(d \log n)$ time for a worst case edge even for $d = \exp(\Theta(\sqrt{\log n}))$.

Jakab Tardos, Michael Kapralov
EPFL
jakab.tardos@epfl.ch, michael.kapralov@epfl.ch

Slobodan Mitrovic
MIT
slobo@mit.edu

Ashkan Norouzi-Fard
Google Zurich
ashkan.afn@gmail.com

CP30

Competitive Online Search Trees on Trees

We consider the design of adaptive data structures for searching elements of a tree-structured space. We use a natural generalization of the rotation-based online binary

search tree model in which the underlying search space is the set of vertices of a tree. This model is based on a simple structure for decomposing graphs, previously known under several names including elimination trees, vertex rankings, and tubings. The model is equivalent to the classical binary search tree model exactly when the underlying tree is a path. We describe an online $O(\log \log n)$-competitive search tree data structure in this model, matching the best known competitive ratio of binary search trees. Our method is inspired by Tango trees, an online binary search tree algorithm, but critically needs several new notions including one which we call Steiner-closed search trees, which may be of independent interest. Moreover our technique is based on a novel use of two levels of decomposition, first from search space to a set of Steiner-closed trees, and secondly from these trees into paths.

Grigorios Koumoutsos
Universite libre de Bruxelles (ULB)
gregkoumoutsos@gmail.com

Prosenjit Bose
Carleton University
jit@scs.carleton.ca

Jean Cardinal
ULB, Brussels
jcardin@ulb.ac.be

John Iacono
The Polytechnic Institute of New York University
johniacono@gmail.com

Stefan Langerman
Universite Libre de Bruxelles
stefan.langerman@ulb.ac.be

CP30

Online Scheduling Via Learned Weights

Online algorithms are a hallmark of worst case optimization under uncertainty. On the other hand, in practice, the input is often far from worst case, and has some predictable characteristics. A recent line of work has shown how to use machine learned predictions to circumvent lower bounds on competitive ratios in classic online problems such as ski rental and caching. We study how predictive techniques can be used to break through worst case barriers in online scheduling. The makespan minimization problem with restricted assignments is a classic problem in online scheduling theory. Worst case analysis gives $\Omega(\log m)$ lower bounds on the competitive ratio in the online setting. We identify a robust quantity that can be predicted and then used to guide online algorithms to achieve better performance. Our predictions have dimension linear in the number of machines and can be learned using standard methods. The performance guarantees of our algorithms depend on the accuracy of the predictions. Given predictions with error η, we show how to construct $O(\log \eta)$ competitive fractional assignments. We then give an $O((\log \log m)^3)$-competitive online algorithm that rounds any fractional assignment into an integral schedule. Altogether, we give algorithms that, equipped with predictions with error η, achieve $O(\log \eta\ (\log \log m)^3)$ competitive ratios, breaking the $\Omega(\log m)$ lower bound even for moderately accurate predictions.

Silvio Lattanzi
Google Zurich
silviol@google.com

Thomas J. Lavastida, Benjamin Moseley
Carnegie Mellon University
tlavasti@andrew.cmu.edu, moseleyb@andrew.cmu.edu

Sergei Vassilvitskii
Google Research
sergeiv@google.com

CP30

Quantifying the Burden of Exploration and the Unfairness of Free Riding

We consider the multi-armed bandit setting with a twist. Rather than having just one decision maker deciding which arm to pull in each round, we have n different decision makers (agents). In the simple stochastic setting, we show that a “free-riding’ agent observing another “self-reliant’ agent can achieve just $O(1)$ regret, as opposed to the regret lower bound of $\Omega(\log t)$ when one decision maker is playing in isolation. This result holds whenever the self-reliant agent’s strategy satisfies either one of two assumptions: (1) each arm is pulled at least $\gamma \ln t$ times in expectation for a constant γ that we compute, or (2) the self-reliant agent achieves $o(t)$ realized regret with high probability. Both of these assumptions are satisfied by standard zero-regret algorithms. Under the second assumption, we further show that the free rider only needs to observe the number of times each arm is pulled by the self-reliant agent, and not the rewards realized. We show that the free rider can achieve $O(1)$ regret in a linear contextual setting whenever the free rider’s context is a small (in L_2-norm) linear combination of other agents’ contexts and all other agents pull each arm $\Omega(\log t)$ times with high probability. Again, this condition on the self-reliant players is satisfied by standard zero-regret algorithms like UCB. We also prove a number of lower bounds.

Christopher Jung, Sampath Kannan
University of Pennsylvania
chrjung@seas.upenn.edu, kannan@cis.upenn.edu

Neil Lutz
Iowa State University
nlutz@iastate.edu

CP30

Interleaved Caching with Access Graphs

We consider a semi-online model for caching in which request sequences are generated by walks on a directed graph, called the access graph. The caching algorithm knows the access graph but not the actual request sequences. We then extend this model to multiple access graphs, where request sequences from the access graphs are interleaved arbitrarily and presented to the caching algorithm. For both these problems, we obtain tight upper and lower bounds on the competitive ratio; our bounds depend on a structural property of the access graph. Our work is motivated by multitasking systems with shared cache, where each task can be abstracted as a directed graph with nodes corresponding to data access and directed edges corresponding to the control flow of the task.

Ravi Kumar
Google Inc.
ravi.k53@gmail.com

Manish Purohit
Department of Computer Science
University of Maryland, College Park.
purohitmanish89@gmail.com

Zoya Svitkina, Erik Vee
Google
Mountain View, CA
zoya@google.com, erikvee@google.com

CP30

Near-Optimal Bounds for Online Caching with Machine Learned Advice

In the model of online caching with machine learned advice, introduced by Lykouris and Vassilvitskii, the goal is to solve the caching problem with an online algorithm that has access to next-arrival predictions: when each input element arrives, the algorithm is given a prediction of the next time when the element will reappear. The traditional model for online caching suffers from an $\Omega(\log k)$ competitive ratio lower bound (on a cache of size k). In contrast, the augmented model admits algorithms which beat this lower bound when the predictions have low error, and asymptotically match the lower bound when the predictions have high error, even if the algorithms are oblivious to the prediction error. In particular, Lykouris and Vassilvitskii showed that there is a prediction-augmented caching algorithm with a competitive ratio of $O(1 + \min(\sqrt{\eta/\textsc{opt}}, \log k))$ when the overall ℓ_1 prediction error is bounded by η, and $\textsc{opt}$ is the cost of the optimal offline algorithm. The dependence on k in the competitive ratio is optimal, but the dependence on $\eta/\textsc{opt}$ may be far from optimal. In this work, we make progress towards closing this gap. Our contributions are twofold. First, we provide an improved algorithm with a competitive ratio of $O(1+\min((\eta/\textsc{opt})/k, 1)\log k)$. Second, we provide a lower bound of $\Omega(\log \min((\eta/\textsc{opt})/(k \log k), k))$.

Dhruv W. Rohatgi
Massachusetts Institute of Technology
drohatgi@mit.edu

CP31

Diameter Computation on H-Minor Free Graphs and Graphs of Bounded (distance) Vc-Dimension

There are several graph classes for which we can compute the diameter in truly subquadratic time such as bounded-treewidth graphs, interval graphs and planar graphs. We propose to study unweighted graphs of constant *distance VC-dimension* as a broad generalization of many such classes – where the distance VC-dimension of a graph G is defined as the VC-dimension of its ball hypergraph.

- Our first result is that on graphs of distance VC-dimension at most d, we can decide if the diameter is larger than k in time $\tilde{\mathcal{O}}(k \cdot mn^{1-\varepsilon_d})$, where $\varepsilon_d \in (0;1)$ only depends on d.
- Then, we get the first truly subquadratic-time algorithm for *constant* diameter computation on all the *nowhere dense* graph classes. The latter classes include all proper minor-closed graph classes, bounded-degree graphs and graphs of bounded expansion.
- Finally, we show how to remove the dependency on k in some particular cases, and in particular for *any* graph class that excludes a fixed graph H as a minor. As a result for all such graphs one obtains a truly subquadratic-time algorithm for computing their diameter.

Our approach is based on the work of Chazelle and Welzl who proved the existence of spanning paths with strongly sublinear *stabbing number* for every hypergraph of constant VC-dimension.

Guillaume Ducoffe
ICI National Institute for Research and Development in Informatics
guillaume.ducoffe@ici.ro

Michel Habib
IRIF
habib@irif.fr

Laurent Viennot
Inria
laurent.viennot@inria.fr

CP31

A Face Cover Perspective to ℓ_1 Embeddings of Planar Graphs

It was conjectured by Gupta et al. [Combinatorica04] that every planar graph can be embedded into ℓ_1 with constant distortion. However, given an n-vertex weighted planar graph, the best upper bound on the distortion is only $O(\sqrt{\log n})$, by Rao [SoCG99]. In this paper we study the case where there is a set K of terminals, and the goal is to embed only the terminals into ℓ_1 with low distortion. In a seminal paper, Okamura and Seymour [J.Comb.Theory81] showed that if all the terminals lie on a single face, they can be embedded isometrically into ℓ_1. The more general case, where the set of terminals can be covered by γ faces, was studied by Lee and Sidiropoulos [STOC09] and Chekuri et al. [J.Comb.Theory13]. The state of the art is an upper bound of $O(\log \gamma)$ by Krauthgamer, Lee and Rika [SODA19]. Our contribution is a further improvement on the upper bound to $O(\sqrt{\log \gamma})$. Since every planar graph has at most $O(n)$ faces, any further improvement on this result, will be a major breakthrough, directly improving upon Rao's long standing upper bound. Moreover, it is well known that the flow-cut gap equals to the distortion of the best embedding into ℓ_1. Therefore, our result provides a polynomial time $O(\sqrt{\log \gamma})$-approximation to the sparsest cut problem on planar graphs, for the case where all the demand pairs can be covered by γ faces.

Arnold Filtser
Columbia University
arnold273@gmail.com

CP31

A Blossom Algorithm for Maximum Edge-Disjoint T-Paths

Let $G = (V, E)$ be a multigraph with a set $T \subseteq V$ of terminals. A path in G is called a T-path if its ends are distinct vertices in T and no internal vertices belong to T. In 1978, Mader showed a characterization of the maximum number of edge-disjoint T-paths. The original proof was not constructive, and hence it did not suggest an efficient algorithm. In this paper, we provide a combinatorial, deterministic algorithm for finding the maximum number of edge-disjoint T-paths. The algorithm adopts an augmenting path approach. More specifically, we introduce a novel concept of augmenting walks in auxiliary labeled

graphs to capture a possible augmentation of the number of edge-disjoint T-paths. To design a search procedure for an augmenting walk, we introduce blossoms analogously to the blossom algorithm of Edmonds (1965) for the matching problem, while it is neither a special case nor a generalization of the present problem. When the search procedure terminates without finding an augmenting walk, the algorithm provides a certificate for the optimality of the current edge-disjoint T-paths. Thus the correctness argument of the algorithm serves as an alternative direct proof of Mader's theorem on edge-disjoint T-paths. The algorithm runs in $O(|V| \cdot |E|^2)$ time, which is much faster than the best known deterministic algorithm based on a reduction to the linear matroid parity problem.

Satoru Iwata, Yu Yokoi
University of Tokyo
iwata@mist.i.u-tokyo.ac.jp, yokoi@nii.ac.jp

CP31

Multi-Transversals for Triangles and the Tuza Conjecture

In this paper, we study a primal and dual relationship about triangles: For any graph G, let $\nu(G)$ be the maximum number of edge-disjoint triangles in G, and $\tau(G)$ be the minimum subset F of edges such that $G \setminus F$ is triangle-free. It is easy to see that $\nu(G) \leq \tau(G) \leq 3\nu(G)$, and in fact, this rather obvious inequality holds for a much more general primal-dual relation between k-hyper matching and covering in hypergraphs. Tuza conjectured in 1990 that $\tau(G) \leq 2\nu(G)$, and this question has received attention from various groups of researchers in discrete mathematics, settling various special cases such as planar graphs and generalized to 6-degenerate graphs, some cases of minor-free graphs, and very dense graphs. In this paper, we provide a proof of a non-trivial consequence of the conjecture; that is, for every $k \geq 2$, there exist a (multi)-set $F \subseteq E(G) : |F| \leq 2k\nu(G)$ such that each triangle in G overlaps at least k elements in F. Our result can be seen as a strengthened statement of Krivelevich's result on the fractional Tuza conjecture (and we give some examples illustrating this.) The main technical ingredient of our result is a charging argument, that locally identifies edges in F (to remove) based on a local view of the packing solution. This idea might be useful in further studying the primal-dual relations in general and the Tuza conjecture in particular.

Parinya Chalermsook
Aalto University, Espoo, Finland
parinya.chalermsook@aalto.fi

Samir Khuller, Pattara Sukprasert
Northwestern University, Evanston, USA
samir.khuller@northwestern.edu,
pattara.sk127@gmail.com

Sumedha Uniyal
Aalto University, Espoo, Finland
sumedha.uniyal@aalto.fi

CP31

A Strongly Polynomial Algorithm for Finding a Shortest Non-Zero Path in Group-Labeled Graphs

We study the problem of finding a shortest non-zero path/cycle in group-labeled graphs with nonnegative edge length, which includes the following two types of tractable variants in undirected graphs, depending on the group in question. One is the parity-constrained shortest path/cycle problem, and the other is computing a shortest noncontractible cycle in surface-embedded graphs. For the shortest non-zero path problem with respect to finite abelian groups, Kobayashi and Toyooka (2017) proposed a randomized, pseudopolynomial algorithm via permanent computation. For a slightly more general class of groups,Yamaguchi (2016) showed a reduction of the problem to the weighted linear matroid parity problem. In particular, some cases are solved in strongly polynomial time with the aid of a deterministic, polynomial algorithm for the weighted linear matroid parity problem developed by Iwata and Kobayashi (2017), which generalizes a well-known fact that the shortest odd/even path problem is solved via weighted matching. In this paper, as the first general solution independent of the group, we present a rather simple, deterministic, and strongly polynomial algorithm for the shortest non-zero path problem. The algorithm is based on Dijkstra's algorithm for the unconstrained shortest path problem and Edmonds' blossom shrinking technique in matching algorithms, and clarifies a common tractable feature behind the parity and topological constraints in the shortest path/cycle problem.

Yutaro Yamaguchi
Osaka University
Department of Information and Physical Sciences
yutaro_yamaguchi@ist.osaka-u.ac.jp

CP32

Domain Reduction for Monotonicity Testing: A $o(d)$ Tester for Boolean Functions in d-Dimensions

We describe a $\tilde{O}(d^{5/6})$-query monotonicity tester for Boolean functions $f : [n]^d \to \{0, 1\}$ on the n-hypergrid. This is the first $o(d)$ tester with query complexity independent of n. Motivated by this independence of n, we initiate the study of monotonicity testing of measurable Boolean functions $f : \mathbf{R}^d \to \{0, 1\}$, where the distance is measured with respect to a product distribution over $\mathbf{R}^d$. We give a $\tilde{O}(d^{5/6})$-query monotonicity tester for such functions. Our main technical result is a domain reduction theorem for monotonicity. For any function $f : [n]^d \to \{0, 1\}$, let ϵ_f be its distance to monotonicity. Consider the restriction $\hat{f}$ of the function on a random $[k]^d$ sub-hypergrid of the original domain. We show that for $k = \mathrm{poly}(d/\epsilon_f)$, the expected distance of the restriction is $\mathbf{E}[\epsilon_{\hat{f}}] = \Omega(\epsilon_f)$. Previously, such a result was only known for $d = 1$ (Berman-Raskhodnikova-Yaroslavtsev, STOC 2014). Our result for testing Boolean functions over $[n]^d$ then follows by applying the $O(d^{5/6})\mathrm{poly}(1/\epsilon, \log n, \log d)$-query hypergrid tester of Black-Chakrabarty-Seshadhri (SODA 2018). To obtain the tester over $\mathbf{R}^d$, we use standard measure theoretic tools to reduce monotonicity testing of a measurable function f to monotonicity testing of a discretized version of f over a hypergrid domain $[N]^d$ for large, but finite, N.

Hadley Black
University of California, Santa Cruz
hablack@cs.ucla.edu

Deeparnab Chakrabarty
Dartmouth College
deeparnab@gmail.com

C. Seshadhri
University of California, Santa Cruz

sesh@ucsc.edu

CP32

Testing Convexity of Functions over Finite Domains

We establish new upper and lower bounds on the number of queries required to test convexity of functions over various discrete domains.

1. We provide a simplified version of the non-adaptive convexity tester on the line. We re-prove the upper bound $O(\frac{\log(\epsilon n)}{\epsilon})$ in the usual uniform model, and prove an $O(\frac{\log n}{\epsilon})$ upper bound in the distribution-free setting.

2. We show a tight lower bound of $\Omega(\frac{\log(\epsilon n)}{\epsilon})$ queries for testing convexity of functions $f : [n] \to \mathbb{R}$ on the line. This lower bound applies to both adaptive and non-adaptive algorithms, and matches the upper bound from item 1, showing that adaptivity does not help in this setting.

3. Moving to higher dimensions, we consider the case of a stripe $[3] \times [n]$. We construct an *adaptive* tester for convexity of functions $f \colon [3] \times [n] \to \mathbb{R}$ with query complexity $O(\log^2 n)$. We also show that any *non-adaptive* tester must use $\Omega(\sqrt{n})$ queries in this setting. Thus, adaptivity yields an exponential improvement for this problem.

4. For functions $f \colon [n]^d \to \mathbb{R}$ over domains of dimension $d \geq 2$, we show a non-adaptive query lower bound $\Omega((\frac{n}{d})^{\frac{d}{2}})$.

Eric Blais
University of Waterloo
eric.blais@uwaterloo.ca

Aleksandrs Belovs
University of Latvia
aleksandrs.belovs@lu.lv

Abhinav Bommireddi
University of Waterloo
vabommir@uwaterloo.ca

CP32

Reconstruction under Outliers for Fourier-Sparse Functions

We consider the problem of learning an unknown f with a sparse Fourier spectrum in the presence of outlier noise. In particular, the algorithm has access to a noisy oracle for (an unknown) f such that (1) the Fourier spectrum of f is k-sparse; (2) at any query point x, the oracle returns y such that with probability $1 - \rho, |y - f(x)| < \epsilon$. However, with probability ρ, the error y-f(x) can be arbitrarily large. We study Fourier sparse functions over both the discrete cube $\{0,1\}^n$ and the torus $[0,1)$ and for both these domains, we design efficient algorithms which can tolerate any $\rho < 1/2$ fraction of outliers. We note that the analogous problem for low-degree polynomials has recently been studied in several works [AK03,GZ16,KKP17] and similar algorithmic guarantees are known in that setting. While our main results pertain to the case where the location of the outliers, i.e., x such that $|y - f(x)\| > \epsilon$ is randomly distributed, we also study the case where the outliers are adversarially located. In particular, we show that over the torus, assuming that the Fourier transform satisfies a certain granularity condition, there is a sample efficient algorithm to tolerate $\rho =$ Omega(1) fraction of outliers and further, that this is not possible without such a granularity condition. Finally, while not the principal thrust, our techniques also allow us non-trivially improve on learning low-degree functions f on the hypercube in the presence of adversarial outlier noise.

Xue Chen
Northwestern University
xue.chen1@northwestern.edu

Anindya De
University of Pennsylvania
anindyad@seas.upenn.edu

CP32

Approximating the Distance to Monotonicity of Boolean Functions

We design a nonadaptive algorithm that, given a Boolean function $f : \{0,1\}^n \to \{0,1\}$ which is α-far from monotone, makes poly$(n, 1/\alpha)$ queries and returns an $\tilde{O}(\sqrt{n})$-approximation to the distance of f to monotonicity. Furthermore, we show that for any constant $\kappa > 0$, approximating the distance to monotonicity up to $n^{1/2-\kappa}$-factor requires 2^{n^κ} nonadaptive queries, thereby ruling out a poly$(n, 1/\alpha)$-query nonadaptive algorithm for such approximations. Approximating the distance to a property is equivalent to tolerantly testing that property. Our lower bound stands in contrast to standard (non-tolerant) testing of monotonicity that can be done nonadaptively with $\tilde{O}(\sqrt{n}/\epsilon^2)$ queries. We obtain our lower bound by proving an analogous bound for erasure-resilient testers. An α-erasure-resilient tester for a desired property gets oracle access to a function that has at most an α fraction of values erased. The tester has to accept (with probability at least 2/3) if the erasures can be filled in to ensure that the resulting function has the property and to reject (with probability at least 2/3) if every completion of erasures results in a function that is ϵ-far from having the property. Our method yields the same lower bounds for unateness and being a k-junta. These lower bounds improve exponentially on the existing lower bounds for these properties.

Erik Waingarten
Columbia University
eaw@cs.columbia.edu

Ramesh Krishnan S. Pallavoor, Sofya Raskhodnikova
Boston University
rameshkp@bu.edu, sofya@bu.edu

CP32

Computing and Testing Small Connectivity in Near-Linear Time and Queries Via Fast Local Cut Algorithms

Consider the following "local' cut-detection problem in a directed graph: We are given a starting vertex s and need to remove at most k edges so that at most ν edges can be reached from s (a local cut) or output $\perp$ to indicate that no such cut exists. If we are given query accesses to the input graph, then this problem can in principle be solved without reading the whole graph and with query complexity depending on k and ν. In this paper we consider a slack variant of this problem where, when such cut exists, we can output a cut with up to $O(k\nu)$ edges reachable from s. We present a simple randomized algorithm with $O(k^2\nu)$ time and $O(k\nu)$ queries for the above variant. We demonstrate that these local algorithms are versatile primitives for designing substantially improved algorithms for classic graph problems by providing the following three applications. 1. A randomized algorithm for the classic k-vertex connectivity problem that takes near-linear time when $k = O(\text{polylog}(n))$, namely $\tilde{O}(m + nk^3)$ time. This improves the recent algorithm with $\tilde{O}(m + n^{4/3}k^{7/3})$ time [Nanongkai et al., STOC'19]. 2. Property testing algorithms for k-edge and -vertex connectivity with query complexities that are near-linear in k, exponentially improving the state-of-the-art. 3. A faster algorithm for computing the maximal k-edge connected subgraphs, improving prior work of Chechik et al. [SODA'17].

Sebastian Forster
University of Salzburg
forster@cs.sbg.ac.at

Danupon Nanongkai
KTH Royal Institute of Technology
Sweden
danupon@kth.se

Thatchaphol Saranurak
Toyota Technological Institute at Chicago
saranurak@ttic.edu

Liu Yang
Independent Researcher)
liu.yang0900@outlook.com

Sorrachai Yingchareonthawornchai
Aalto University
sorrachai.cp@gmail.com

CP33

The Two-Sided Game of Googol and Sample-Based Prophet Inequalities

The secretary problem or the game of Googol are classic models for online selection problems that have received significant attention in the last five decades. We consider a variant of the problem and explore its connections to data-driven online selection. Specifically, we are given n cards with arbitrary non-negative numbers written on both sides. The cards are randomly placed on n consecutive positions on a table, and for each card, the visible side is also selected at random. The player sees the visible side of all cards and wants to select the card with the maximum hidden value. To this end, the player flips the first card, sees its hidden value and decides whether to pick it or drop it and continue with the next card. We study algorithms for two natural objectives. In the first one, as in the secretary problem, the player wants to maximize the probability of selecting the maximum hidden value. We show that this can be done with probability at least 0.45292. In the second one, similar to the prophet inequality, the player maximizes the expectation of the selected hidden value. We show a guarantee of at least 0.63518 with respect to the expected maximum hidden value. We apply our results to the prophet secretary problem with unknown distributions, but with access to a single sample from each distribution. Our guarantee improves upon $1 - 1/e$ for this problem, which is the currently best known guarantee and only works for the i.i.d. case.

José Correa, Andrés Cristi, Boris Epstein, José Soto
Universidad de Chile
correa@uchile.cl, andres.cristi@ing.uchile.cl, boris.epstein@ing.uchile.cl, jsoto@dim.uchile.cl

CP33

A Truthful Cardinal Mechanism for One-Sided Matching

We revisit the well-studied problem of designing mechanisms for one-sided matching markets, where a set of n agents needs to be matched to a set of n heterogeneous items. Each agent i has a value $v_{i,j}$ for each item j, and these values are private information that the agents may misreport if doing so leads to a preferred outcome. Ensuring that the agents have no incentive to misreport requires a careful design of the matching mechanism, and mechanisms proposed in the literature mitigate this issue by eliciting only the *ordinal* preferences of the agents, i.e., their ranking of the items from most to least preferred. However, the efficiency guarantees of these mechanisms are based only on weak measures that are oblivious to the underlying values. In this paper we achieve stronger performance guarantees by introducing a mechanism that truthfully elicits the full *cardinal* preferences of the agents, i.e., all of the $v_{i,j}$ values. We evaluate the performance of this mechanism using the much more demanding Nash bargaining solution as a benchmark, and we prove that our mechanism significantly outperforms all ordinal mechanisms (even non-truthful ones). To prove our approximation bounds, we also study the population monotonicity of the Nash bargaining solution in the context of matching markets, providing both upper and lower bounds which are of independent interest.

Vasilis Gkatzelis
Drexel University, Computer Science department
gkatz@drexel.edu

Rediet Abebe
Cornell University
Harvard University
rabebe@fas.harvard.edu

Richard Cole
NYU
cole@cs.nyu.edu

Jason Hartline
Northwestern University
hartline@northwestern.edu

CP33

Instance-Optimality in the Noisy Value-and Comparison-Model

Motivated by crowdsourced computation, peer-grading, and recommendation systems, Braverman, Mao and Weinberg studied the *query* and *round* complexity of fundamental problems such as finding the maximum (MAX), finding all elements above a certain value (THRESHOLD-v) or computing the top$-k$ elements (TOP-k) in a noisy environment. For example, consider the task of selecting papers for a conference. This task is challenging due to the crowdsourcing nature of peer reviews: the results of reviews are noisy and it is necessary to parallelize the review process as much as possible. We study the noisy value model and the noisy comparison model: In the *noisy value model*, a reviewer is asked to evaluate a single element: "What is the value of paper i?' (accept). In the *noisy comparison model* (introduced in the seminal work of Feige, Peleg, Raghavan and Upfal) a reviewer is asked to do a pairwise comparison: "Is paper i better than paper j?' In this paper, we introduce new lower bound techniques for these classic problems. We complement these results with simple algorithms which show that our lower bounds are almost tight. We then go beyond the worst-case and address the question of the importance of knowledge of the instance by providing, for a large range of parameters, instance-optimal algorithms with respect to the query complexity.

Frederik Mallmann-Trenn
King's College London
frederik.mallmann-trenn@kcl.ac.uk

Claire Mathieu
CNRS, IRIF
clairemmathieu@gmail.com

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

CP33

Competitive Analysis with a Sample and the Secretary Problem

We extend the standard online worst-case model to accommodate past experience which is available to the online player in many practical scenarios. We do this by revealing a random sample of the adversarial input to the online player ahead of time. The online player competes with the expected optimal value on the part of the input that arrives online. Our model bridges between existing online stochastic models (e.g., items are drawn i.i.d. from a distribution) and the online worst-case model. We also extend in a similar manner (by revealing a sample) the online random-order model. We study the classical secretary problem in our new models. In the worst-case model we present a simple online algorithm with optimal competitive-ratio for any sample size. In the random-order model, we also give a simple online algorithm with an almost tight competitive-ratio for small sample sizes. Interestingly, we prove that for a large enough sample, no algorithm can be simultaneously optimal both in the worst-cast and random-order models.

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

David Naori, Danny Raz
The Technion Israel Institute of Technology
dnaori@cs.technion.ac.il, danny@cs.technion.ac.il

CP33

Cake Cutting on Graphs: A Discrete and Bounded Proportional Protocol

The classical cake cutting problem studies how to find fair allocations of a heterogeneous and divisible resource among multiple agents. Two of the most commonly studied fairness concepts in cake cutting are proportionality and envy-freeness. It is well known that a proportional allocation among n agents can be found efficiently via simple protocols. For envy-freeness, in a recent breakthrough, Aziz and Mackenzie proposed a discrete and bounded envy-free protocol for any number of players. However, the protocol suffers from high multiple-exponential query complexity and it remains open to find simpler and more efficient envy-free protocols. In this paper we consider a variation of the cake cutting problem by assuming an underlying graph over the agents whose edges describe their acquaintance relationships, and agents evaluate their share relatively to those of their neighbors. An allocation is called locally proportional if each agent thinks she receives at least the average value over her neighbors. Local proportionality is implied by envy-freeness, but no simple protocol is known to produce a locally proportional allocation for general graphs. In this paper we answer this question by presenting a discrete and bounded locally proportional protocol for any given graph. Our protocol has a query complexity of only single exponential, which is significantly smaller than the six towers of n query complexity of the envy-free protocol given by Aziz and Mackenzie.

Zhijie Zhang
Institute of Computing Technology
Chinese Academy of Sciences
zhangzhijie@ict.ac.cn

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

Xiaoming Sun
Institute of Computing Technology
Chinese Academy of Sciences
sunxiaoming@ict.ac.cn

Hao Wu
University of Chinese Academy of Sciences
wuhao164@mails.ucas.ac.cn

Jialin Zhang
Institute of Computing Technology
Chinese Academy of Sciences
zhangjialin@ict.ac.cn

Wei Zi
University of Chinese Academy of Sciences
ziwei16@mails.ucas.ac.cn

CP34

Reconstruction of Depth-4 Multilinear Circuits

We present a deterministic algorithm for reconstructing multilinear $\Sigma\Pi\Sigma\Pi(k)$ circuits, i.e. multilinear depth-4 circuits with fan-in k at the top + gate. For any fixed k, given black-box access to a polynomial $f \in F[x_1, x_2, \ldots, x_n]$ computable by a multilinear $\Sigma\Pi\Sigma\Pi(k)$ circuit of size s, the algorithm runs in time quasi-poly$(n, s, |F|)$ and outputs a multilinear $\Sigma\Pi\Sigma\Pi(k)$ circuit of size quasi-poly(n, s) that computes f. Our result solves an open problem posed in [GKL12] (STOC, 2012). Indeed, prior to our work, efficient reconstruction algorithms for multilinear $\Sigma\Pi\Sigma\Pi(k)$ circuits were known only for the case of $k = 2$ [GKL12, Volkovich17].

Vishwas Bhargava
Rutgers University
vishwas1384@gmail.com

Shubhangi Saraf
Department of Mathematics and Department of

Computer Science
Rutgers University
shubhangi.saraf@gmail.com

Ilya Volkovich
Department of EECS, CSE Division
University of Michigan
ilyavol@umich.edu

CP34

Approximately Counting and Sampling Small Witnesses Using a Colourful Decision Oracle

In this paper, we prove “black box’ results for turning algorithms which decide whether or not a witness exists into algorithms to approximately count the number of witnesses, or to sample from the set of witnesses approximately uniformly, with essentially the same running time. We do so by extending the framework of Dell and Lapinskas (STOC 2018), which covers decision problems that can be expressed as edge detection in bipartite graphs given limited oracle access; our framework covers problems which can be expressed as edge detection in arbitrary k-hypergraphs given limited oracle access. (Simulating this oracle generally corresponds to invoking a decision algorithm.) This includes many key problems in both the fine-grained setting (such as k-SUM, k-OV and weighted k-Clique) and the parameterised setting (such as induced subgraphs of size k or weight-k solutions to CSPs). From an algorithmic standpoint, our results will make the development of new approximate counting algorithms substantially easier; indeed, it already yields a new state-of-the-art algorithm for approximately counting graph motifs, improving on Jerrum and Meeks (JCSS 2015) unless the input graph is very dense and the desired motif very small. Our k-hypergraph reduction framework generalises and strengthens results in the graph oracle literature due to Beame et al. (ITCS 2018) and Bhattacharya et al. (CoRR abs/1808.00691).

Holger Dell
ITU Copenhagen
hold@itu.dk

John Lapinskas
University of Birmingham
john.lapinskas@bristol.ac.uk

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

CP34

Parameterized Complexity and Approximability of Directed Odd Cycle Transversal

A *directed odd cycle transversal* of a directed graph (digraph) D is a vertex set S that intersects every *odd directed cycle* of D. In the Directed Odd Cycle Transversal (DOCT) problem, the input consists of a digraph D and an integer k. The objective is to determine whether there exists a directed odd cycle transversal of D of size at most k. In this paper, we settle the parameterized complexity of DOCT when parameterized by the solution size k by showing that DOCT does not admit an algorithm with running time $f(k)n^{O(1)}$ unless $FPT = W[1]$. On the positive side, we give a factor 2 fixed-parameter approximation (FPT approximation) algorithm for the problem. More precisely, our algorithm takes as input D and k, runs in time $2^{O(k^2)}n^{O(1)}$, and either concludes that D does not have a directed odd cycle transversal of size at most k, or produces a solution of size at most $2k$. Finally, assuming gap-ETH, we show that there exists an $\epsilon > 0$ such that DOCT does not admit a factor $(1 + \epsilon)$ FPT-approximation algorithm. In fact, our lower bound holds even under the weaker hypothesis that for some $\epsilon > 0$, there is no $f(k)n^{O(1)}$ time algorithm that distinguishes between a satisfiable Binary CSP instance from one where every assignment violates at least an ϵ fraction of the constraints.

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

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Meirav Zehavi
BGU
meiravze@bgu.ac.il

Ramanujan M. S.
University of Warwick
r.maadapuzhi-sridharan@warwick.ac.uk

CP34

Improved Bounds for Centered Colorings

A vertex coloring ϕ of a graph G is *p-centered* if for every connected subgraph H of G either ϕ uses more than p colors on H or there is a color that appears exactly once on H. Centered colorings allow to capture notions of sparsity of graphs: A class of graphs has bounded expansion if and only if there is a function f such that for every p, every graph in the class admits a p-centered coloring using at most $f(p)$ colors. We show upper bounds for the maximum number of colors needed in a p-centered coloring of graphs from several widely studied graph classes: (1) planar graphs admit p-centered colorings with $O(p^3 \log p)$ colors; (2) bounded degree graphs admit p-centered colorings with $O(p)$ colors while it was conjectured that they may require exponential number of colors in p; (3) graphs avoiding a fixed graph as a topological minor admit p-centered colorings with a polynomial in p number of colors. All these upper bounds imply polynomial time algorithms. Prior to this work there were no non-trivial lower bounds known. We show that: (4) there are graphs of treewidth t that require $\binom{p+t}{t}$ colors in any p-centered coloring and this bound matches the upper bound; (5) there are planar graphs that require $\Omega(p^2 \log p)$ colors in any p-centered coloring. We also give asymptotically tight bounds for outerplanar graphs and planar graphs of treewidth 3.

Michal Debski
Warsaw University of Technology
Masaryk University, Brno
michal.debski87@gmail.com

Stefan Felsner
Technical University of Berlin
Institute for Mathematics
felsner@math.tu-berlin.de

Piotr Micek
Jagellionian University
piotr.Micek@tcs.uj.edu.pl

Felix Schröder
Technische Universit
fschroed@math.tu-berlin.de

CP34

Counting and Finding Homomorphisms Is Universal for Parameterized Complexity Theory

Counting homomorphisms from a graph H into another graph G is a fundamental problem of (parameterized) counting complexity theory. In this work, we study the case where *both* graphs H and G stem from given classes of graphs: $H \in \mathcal{H}$ and $G \in \mathcal{G}$. By this, we combine the structurally restricted version of this problem, with the language-restricted version. Our main result is a construction based on Kneser graphs that associates every problem $\mathtt{P}$ in $\#\mathsf{W}[1]$ with two classes of graphs $\mathcal{H}$ and $\mathcal{G}$ such that the problem $\mathtt{P}$ is *equivalent* to the problem $\#\mathtt{HOM}(\mathcal{H} \to \mathcal{G})$ of counting homomorphisms from a graph in $\mathcal{H}$ to a graph in $\mathcal{G}$. In view of Ladner's seminal work on the existence of NP-intermediate problems [J.ACM'75] and its adaptations to the parameterized setting, a classification of the class $\#\mathsf{W}[1]$ in fixed-parameter tractable and $\#\mathsf{W}[1]$-complete cases is unlikely. Hence, obtaining a complete classification for the problem $\#\mathtt{HOM}(\mathcal{H} \to \mathcal{G})$ seems unlikely. Further, our proofs easily adapt to $\mathsf{W}[1]$. In search of complexity dichotomies, we hence turn to special graph classes. Those classes include line graphs, claw-free graphs, perfect graphs, and combinations thereof, and F-colorable graphs for fixed graphs F.

Marc Roth
Merton College
Oxford University, United Kingdom
marc.roth@merton.ox.ac.uk

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

CP35

Learning Software Constraints Via Installation Attempts

Modern software systems are expected to be secure and contain all the latest features, even when new versions are released multiple times an hour. Each system may include many interacting packages. The problem of installing multiple dependent packages has been extensively studied in the past, yielding some promising solutions that work well in practice. However, these assume that the developers declare all the dependencies and conflicts between the packages. Often, the entire repository structure may not be known up front, for example, when packages are developed by different vendors. In this paper, we present algorithms for learning dependencies, conflicts, and defective packages from installation attempts. Our algorithms use combinatorial data structures to generate queries that test installations and discover the entire dependency structure. A query that the algorithms make corresponds to trying to install a subset of packages and getting Boolean feedback on whether all constraints were satisfied in this subset. Our goal is to minimize the query complexity of the algorithms. We prove lower and upper bounds on the number of queries that these algorithms require to make for different settings of the problem.

Ran Ben Basat
Harvard University
ran@seas.harvard.edu

Maayan Goldstein, Itai Segall
Nokia Bell Labs
maayan.goldstein@nokia-bell-labs.com, itai.segall@nokia-bell-labs.com

CP35

Lock-Free Hopscotch Hashing

We present a lock-free version of Hopscotch Hashing. Hopscotch Hashing is an open addressing algorithm originally proposed by Herlihy, Shavit, and Tzafrir, which is known for fast performance and excellent cache locality. The algorithm allows users of the table to skip or jump over irrelevant entries, allowing quick search, insertion, and removal of entries. Unlike traditional linear probing, Hopscotch Hashing is capable of operating under a high load factor, as probe counts remain small. Our lock-free version improves on both speed, cache locality, and progress guarantees of the original, being a chimera of two concurrent hash tables. We compare our data structure to various other lock-free and blocking hashing algorithms and show that its performance is in many cases superior to existing strategies. The proposed lock-free version overcomes some of the drawbacks associated with the original blocking version, leading to a substantial boost in scalability while maintaining attractive features like physical deletion or probe-chain compression.

Robert Kelly
Maynooth University
rob.kelly@cs.nuim.ie

Barak Pearlmutter, Phil Maguire
Maynooth University, Ireland
barak@pearlmutter.net, pmaguire@cs.nuim.ie

CP35

Memory-Efficient Performance Monitoring on Programmable Switches with Lean Algorithms

Network performance problems are notoriously difficult to diagnose. Prior profiling systems collect performance statistics by keeping information about each network flow, but maintaining per-flow state is not scalable on resource-constrained NIC and switch hardware. Instead, we propose sketch-based performance monitoring using memory that is sublinear in the number of flows. Existing sketches estimate flow monitoring metrics based on flow sizes. In contrast, performance monitoring typically requires combining information across pairs of packets, such as matching a data packet with its acknowledgment to compute a round-trip time. We define a new class of *lean* algorithms that use memory sublinear in both the size of input data and the number of flows. We then introduce lean algorithms for a set of important statistics, such as identifying flows with high latency, loss, out-of-order, or retransmitted packets. We implement prototypes of our lean algorithms on a commodity programmable switch using the P4 language. Our experiments show that lean algorithms detect ~82% of top 100 problematic flows among real-world packet traces using just 40KB memory.

Zaoxing Liu, Samson Zhou
Carnegie Mellon University
zaoxingliu@gmail.com, samsonzhou@gmail.com

Ori Rottenstreich
Technion
ori.rot@gmail.com

Vladimir Braverman
Johns Hopkins University
vova@cs.jhu.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

CP35

Writeback-Aware Caching

The literature on cache replacement, while extensive, neglects to account for the flow of data *to memory*. Motivated by emerging memory technologies and the increasing importance of memory bandwidth and energy consumption, we study the *Writeback-Aware Caching Problem*. This problem modifies traditional caching problems by explicitly accounting for the cost of writing modified data back to memory. In the offline setting with maximum writeback cost $\omega > 0$, we show that writeback-aware caching is NP-complete and Max-SNP hard. Moreover, we show that the optimal deterministic policy when ignoring writebacks is only $(\omega + 1)$-competitive. These negative results hold even when all items have unit size, unit miss cost, and unit writeback cost ($\omega = 1$). To overcome this, we provide practical algorithms to compute upper and lower bounds for the optimal policy on real traces. In the online setting, we present a deterministic replacement policy called *Writeback-Aware Landlord*, and show that it obtains the optimal competitive ratio. Our results hold even for the most general variant in which data items have variable sizes, variable miss costs, and variable writeback costs. Finally, we perform an experimental study on real-world traces showing that Writeback-Aware Landlord outperforms state-of-the-art cache replacement policies when writebacks are costly, thereby illustrating the practical gains of explicitly accounting for writebacks.

Charles Mcguffey, Nathan Beckmann, Phillip Gibbons
Carnegie Mellon University
cmcguffe@cs.cmu.edu, beckmann@cs.cmu.edu,
gibbons@cs.cmu.edu

Bernhard Haeupler
Carnegie Mellon U.
haeupler@cs.cmu.edu

CP35

Parallel Algorithms for Butterfly Computations

We design a framework called ParButterfly that contains new parallel algorithms for the following problems on butterflies, the smallest non-trivial subgraph in bipartite graphs: global counting, per-vertex counting, per-edge counting, and vertex and edge peeling. The main component of these algorithms is aggregating wedges incident on subsets of vertices, and our framework supports different methods for wedge aggregation, including sorting, hashing, histogramming, and batching. ParButterfly also supports different ways of ranking the vertices to speed up counting, including side ordering, approximate and exact degree ordering, and approximate and exact complement coreness ordering. For counting, ParButterfly also supports both exact computation as well as approximate computation via graph sparsification. We prove strong theoretical guarantees on the work and span of the algorithms in ParButterfly. We perform a comprehensive evaluation of all of the algorithms in ParButterfly on a collection of real-world bipartite graphs using a 48-core machine. Our counting algorithms obtain significant parallel speedup, outperforming the fastest sequential algorithms by up to 13.6x with a self-relative speedup of up to 38.5x. Compared to general subgraph counting solutions, we are orders of magnitude faster. Our peeling algorithms achieve self-relative speedups of up to 10.7x and outperform the fastest sequential baseline by up to several orders of magnitude.

Jessica Shi
Massachusetts Institute of Technology
jeshi@mit.edu

Julian Shun
MIT
jshun@mit.edu

CP36

Approximation Schemes for Capacitated Clustering in Doubling Metrics

Motivated by applications in redistricting, we consider the uniform capacitated k-median and uniform capacitated k-means problems in bounded doubling metrics. We provide the first QPTAS for both problems and the first PTAS for the uniform capacitated k-median problem for points in R^2. This is the first improvement over the bicriteria QPTAS for capacitated k-median in low-dimensional Euclidean space of Arora, Raghavan, Rao [STOC 1998] (1 + epsilon-approximation, 1 + epsilon-capacity violation) and arguably the first polynomial-time approximation algorithm for a non-trivial metric.

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

CP36

Baker Game and Polynomial-Time Approximation Schemes

Baker (1994) devised a technique to obtain approximation schemes for many optimization problems restricted to planar graphs; her technique was later extended to more general graph classes. In particular, using the Baker's technique and the minor structure theorem, Dawar et al. (2006) gave Polynomial-Time Approximation Schemes (PTAS) for all monotone optimization problems expressible in the first-order logic when restricted to a proper minor-closed class of graphs. We define a *Baker game* formalizing the notion of repeated application of Baker's technique interspersed with vertex removal, prove that monotone optimization problems expressible in the first-order logic admit PTAS when restricted to graph classes in which the Baker game can be won in a constant number of rounds, and prove *without* use of the minor structure theorem that all proper minor-closed classes of graphs have this property.

Zdenek Dvorak
Computer Science Institute
Charles University, Prague

rakdver@iuuk.mff.cuni.cz

CP36

A Ptas for Subset Tsp in Minor-Free Graphs

We give the first PTAS for the subset Traveling Salesperson Problem (TSP) in H-minor-free graphs. This resolves a long standing open problem in a long line of work on designing PTASes for TSP in minor-closed families initiated by Grigni, Koutsoupias and Papadimitriou in FOCS'95. The main technical ingredient is the construction of a nearly light subset $(1 + \epsilon)$-spanner for any given edge-weighted H-minor-free graph. At the heart of our construction is a new concept, namely *sparse spanner oracles*. We show that spanner oracles with weak sparsity are both necessary and sufficient to construct light subset spanners, even for general graphs. We also present two new results as applications of our new concept.

- We construct an $(1 + \epsilon)$-spanner with lightness $O(\epsilon^{-d+2})$ for any doubling metric of constant dimension d. This improves the earlier lightness bound $\epsilon^{-O(d)}$ obtained by Borradaile, Le and Wulff-Nilsen.
- We show the existence of an $(1+\epsilon)$-spanner with sublinear lightness for any metric of constant correlation dimension. Previously, no spanner with non-trivial lightness was known.

Hung Le
Oregon State University
hungle@uvic.ca

CP36

Approximation Schemes Via Width/Weight Trade-Offs on Minor-Free Graphs

In this paper, we prove a new scaling lemma for vertex weighted minor free graphs that allows for a smooth trade-off between the weight of a vertex set S and the treewidth of $G - S$. More precisely, we show the following. There exists an algorithm that given an H-minor free graph G, a weight function $w : V(G)\mathbb{Q}^+$ and integers t and s, runs in polynomial time, and outputs a subset $S \subseteq V(G)$ of weight at most $d \log n \frac{opt(G,w,t)}{s}$ such that the treewidth of $G - S$ is at most cst. Here, d and c are fixed constants that depend only on H, and $opt(G, w, t)$ is the (unknown) minimum weight of a subset $U \subseteq V(G)$ such that the treewidth of $G - U$ is at most t. This lemma immediately yields the first polynomial-time approximation schemes (PTASes) for Weighted Treewidth-η Vertex Deletion, for $\eta \geq 2$, on graphs of bounded genus and the first PTAS for Weighted Feedback vertex Set on H-minor free graphs. These results effortlessly generalize to include weighted edge deletion problems, to all Weighted Connected Planar $\mathcal{F}$-Deletion problems, and finally to quasi polynomial time approximation schemes (QPTASes) for all of these problems on H-minor free graphs. For most of these problems even constant factor approximation algorithms, even on planar graphs, were not previously known.

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

Fedor V. Fomin
UiB
fomin@ii.uib.no

Daniel Lokshtanov
UCSB
daniello@ucsb.edu

Meirav Zehavi
BGU
meiravze@bgu.ac.il

CP36

Quasi-Polynomial Time Approximation Schemes for the Maximum Weight Independent Set Problem in H-Free Graphs

In the Maximum Independent Set (MIS) problem we are asked to find a set of pairwise nonadjacent vertices in a given graph with the maximum possible cardinality. In general graphs, this classical problem is known to be NP-hard and hard to approximate within a factor of $n^{1-\varepsilon}$ for any $\varepsilon > 0$. Due to this, investigating the complexity of MIS in various graph classes in hope of finding better tractability results is an active research direction. In H-free graphs, that is, graphs not containing a fixed graph H as an induced subgraph, the problem is known to remain NP-hard and APX-hard whenever H contains a cycle, a vertex of degree at least four, or two vertices of degree at least three in one connected component. For the remaining cases, where every component of H is a path or a subdivided claw, the complexity of MIS remains widely open, with only a handful of polynomial-time solvability results for small graphs H. We prove that for every such "possibly tractable' graph H there exists an algorithm that, given an H-free graph G and an accuracy parameter $\varepsilon > 0$, finds an independent set in G of cardinality within a factor of $(1-\varepsilon)$ of the optimum in time exponential in a polynomial of $\log |V(G)|$ and ε^{-1}. That is, we show that for every graph H for which MIS is not known to be APX-hard in H-free graphs, the problem admits a quasi-polynomial time approximation scheme in this graph class.

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

Maria Chudnovsky
Princeton University
mchudnov@math.princeton.edu

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

Stéphan Thomassé
ENS de Lyon
stephan.thomasse@ens-lyon.fr

CP37

Finding Perfect Matchings in Dense Hypergraphs

We show that for any integers $k \geq 3$ and $c \geq 0$ there is a polynomial-time algorithm, that given any n-vertex k-uniform hypergraph H with minimum codegree at least $n/k-c$, finds either a perfect matching in H or a certificate that no perfect matching exists.

Jie Han
Department of Mathematics
University of Rhode Island

jie_han@uri.edu

Peter Keevash
University of Oxford
Mathematical Institute
keevash@maths.ox.ac.uk

CP37

Worst-Case Polylog Incremental Spqr-Trees: Embeddings, Planarity, and Triconnectivity

We show that every labelled planar graph G can be assigned a canonical embedding $\phi(G)$, such that for any planar G' that differs from G by one edge, the number of local changes to the combinatorial embedding needed to get from $\phi(G)$ to $\phi(G')$ is $O(\log n)$. In contrast, there exist embedded graphs where $\Omega(n)$ changes are necessary to accommodate one inserted edge. We provide a matching lower bound of $\Omega(\log n)$ local changes, and, although our upper bound is worst-case, our lower bound holds in the amortized case as well. Our proof is based on BC trees and SPQR trees, and we develop *pre-split* variants of these for general graphs, based on a novel biased heavy-path decomposition, where the structural changes corresponding to edge insertions and deletions in the underlying graph consist of at most $O(\log n)$ basic operations of a particularly simple form. As a secondary result, we show how to maintain the pre-split trees under edge insertions in the underlying graph deterministically in worst case $O(\log^3 n)$ time. Using this, we obtain deterministic data structures for incremental planarity testing, incremental planar embedding, and incremental triconnectivity, that each have worst case $O(\log^3 n)$ update and query time, answering an open question by La Poutré and Westbrook from 1998.

Eva Rotenberg
Technical University of Denmark
eva@rotenberg.dk

Jacob Holm
University of Copenhagen
University of Copenhagen
jaho@di.ku.dk

CP37

Efficiently List-Edge Coloring Multigraphs Asymptotically Optimally

We give polynomial time algorithms for the seminal results of Kahn, who showed that the Goldberg-Seymour and List-Coloring conjectures for (list-)edge coloring multigraphs hold asymptotically. Kahn's arguments are based on the probabilistic method and are non-constructive. Our key insight is that we can combine sophisticated techniques due to Achlioptas, Iliopoulos and Kolmogorov for the analysis of local search algorithms with correlation decay properties of the probability spaces on matchings used by Kahn in order to construct efficient edge-coloring algorithms.

Fotis Iliopoulos, Alistair Sinclair
University of California, Berkeley
fotis.iliopoulos@berkeley.edu, sinclair@berkeley.edu

CP37

Small Memory Robust Simulation of Client-Server Interactive Protocols over Oblivious Noisy Channels

We revisit the problem of low-memory robust simulation of interactive protocols over noisy channels. Haeupler [FOCS 2014] considered robust simulation of two-party interactive protocols over oblivious, as well as adaptive, noisy channels. Since the simulation does not need to have fixed communication pattern, the achieved communication rates can circumvent the lower bound proved by Kol and Raz [STOC 2013]. However, a drawback of this approach is that each party needs to remember the whole history of the simulated transcript. In this paper, we consider low-memory robust simulation of more general client-server interactive protocols, in which a leader communicates with other members/servers, who do not communicate among themselves; this setting can be applied to information-theoretic multi-server Private Information Retrieval (PIR) schemes. We propose an information-theoretic technique that converts any correct PIR protocol that assumes reliable channels, into a protocol which is both correct and private in the presence of a noisy channel while keeping the space complexity to a minimum. Despite the huge attention that PIR protocols have received in the literature, the existing works assume that the parties communicate using noiseless channels.

T-H. Hubert Chan, Zhibin Liang
The University of Hong Kong
hubert@cs.hku.hk, liangzb@connect.hku.hk

Antigoni Polychroniadou
J.P. Morgan AI Research
antigonipoly@gmail.com

Elaine Shi
Cornell University
runting@gmail.com

CP37

Linear Size Sparsifier and the Geometry of the Operator Norm Ball

The Matrix Spencer Conjecture asks whether given n symmetric matrices in $\mathbb{R}^{n\times n}$ with eigenvalues in [-1,1] one can always find signs so that their signed sum has singular values bounded by $O(\sqrt{n})$. The standard approach in discrepancy requires proving that the convex body of all good fractional signings is large enough. However, this question has remained wide open due to the lack of tools to certify measure lower bounds for rather small non-polyhedral convex sets. A seminal result by Batson, Spielman and Srivastava from 2008 shows that any undirected graph admits a linear size spectral sparsifier. Again, one can define a convex body of all good fractional signings. We can indeed prove that this body is close to most of the Gaussian measure. This implies a discrepancy algorithm due to Rothvoss can be used to sample a linear size sparsifer. In contrast to previous methods, we require only a logarithmic number of sampling phases.

Victor Reis, Thomas Rothvoss
University of Washington
voreis@cs.washington.edu, rothvoss@uw.edu

CP38

Bulow-Klemperer-Style Results for Welfare Maximization in Two-Sided Markets

We consider the problem of welfare (and gains-from-trade) maximization in two-sided markets using simple prior-independent mechanisms. Myerson and Satterthwaite [1983] show that no feasible (IR, truthful, BB) mechanism has welfare as high as the optimal-yet-infeasible VCG, which maximizes welfare but runs a deficit. Moreover, the optimal feasible mechanism needs to be carefully tailored to the Bayesian prior, and is extremely complex, eluding a precise description. We present Bulow-Klemperer-style results to circumvent the above in double-auction markets. We suggest the Buyer Trade Reduction (BTR) mechanism, a variant of McAfee's mechanism, which is feasible and simple (in particular: deterministic, truthful, prior-independent, anonymous). First, when the buyers' and sellers' values are i.i.d., we show that for any such market of any size, BTR with one additional buyer with value sampled from the same distribution has expected welfare at least as high as VCG in the original market. We then move to more general settings where the buyers' values are sampled from one distribution and the sellers' from another, focusing on when the former first-order stochastically dominates the latter. We give upper and lower bounds on the number of buyers that, when added, cause BTR in the augmented market to achieve welfare at least as high as VCG in the original market. Our lower bounds apply to a large class of mechanisms, and all our results extend to adding sellers instead of buyers.

Moshe Babaioff
Microsoft Research
moshe@microsoft.com

Kira Goldner
Columbia University
kgoldner@cs.columbia.edu

Yannai A. Gonczarowski
Microsoft Research
yannai@gonch.name

CP38

Inference from Auction Prices

Econometric inference allows an analyst to back out the values of agents in a mechanism from the rules of the mechanism and bids of the agents. This paper gives an algorithm to solve the problem of inferring the values of agents in a dominant-strategy mechanism from: the social choice function implemented by the mechanism and the per-unit prices paid by the agents (the agent bids are not observed). For single-dimensional agents, this inference problem is a multi-dimensional inversion of the payment identity and is feasible only if the payment identity is uniquely invertible. The inversion is unique for single-unit proportional weights social choice functions (common, for example, in bandwidth allocation); and its inverse can be found efficiently. This inversion is not unique for social choice functions that exhibit complementarities. Of independent interest, we extend a result of Rosen (1965), that the Nash equilbria of "concave games' are unique and pure, to an alternative notion of concavity based on Gale and Nikaido (1965).

Jason Hartline
Northwestern U
hartline@eecs.northwestern.edu

Aleck Johnsen
Northwestern University
aleck.johnsen@gmail.com

Denis Nekipelov
University of Virginia
denis@virginia.edu

Zihe Wang
Shanghai University of Finance and Economics
wang.zihe@mail.shufe.edu.cn

CP38

Dominantly Truthful Multi-Task Peer Prediction, with Constant Number of Tasks

In the setting where participants are asked multiple similar possibly subjective multi-choice questions (e.g. restaurants ratings, movie ratings), a series of peer prediction mechanisms are designed to incentivize honest reports and some of them achieve dominant truthfulness: truth-telling is a dominant strategy and strictly dominates other "non-permutation strategy' with some mild conditions. However, a major issue hinders the applicability of these mechanisms: they require either an infinite number of tasks or a finite number of tasks but only have an approximated dominant truthfulness. The existence of a dominantly truthful multi-task peer prediction mechanism that only requires a finite number of tasks remains to be an open question that may have a negative result, even with full prior knowledge. This paper answers this open question by proposing a novel mechanism, DMI-Mechanism (Determinant based Mutual Information Mechanism), that is dominantly truthful with at least 2C tasks where C is the number of choices (C=2 for binary-choice question). To the best of our knowledge, DMI-Mechanism is the first dominantly truthful mechanism that works for a finite number of tasks, not to say a small constant number of tasks. In addition to incentivizing honest reports, DMI-Mechanism can also be transferred into an information evaluation rule that identifies high-quality information without verification when there are at least 3 participants.

Yuqing Kong
Peking University
yuqilng.kong@pku.edu.cn

CP38

Algorithmic Price Discrimination

We consider a generalization of the third degree price discrimination problem studied in [Dirk Bergemann, Benjamin Brooks, and Stephen Mor-ris. The limits of price discrimination], where an intermediary between the buyer and the seller can design market segments to maximize any linear combination of consumer surplus and seller revenue. Unlike in their work, we assume that the intermediary only has partial information about the buyer's value. We consider three different models of information, with increasing order of difficulty. In the first model, we assume that the intermediary's information allows him to construct a probability distribution of the buyer's value. Next we consider the sample complexity model, where we assume that the intermediary only sees samples from this distribution. Finally, we consider a bandit online learning model, where the intermediary can only observe past purchasing decisions of the buyer, rather than her exact value. For each of these models, we present algorithms to compute optimal or near optimal market segmentation.

Rachel Cummings
Georgia Institute of Technology
rachelc@gatech.edu

Nikhil R. Devanur
Microsoft Research, Redmond
nikdev@microsoft.com

Zhiyi Huang, Xiangning Wang
The University of Hong Kong
zhiyi@cs.hku.hk, xnwang@cs.hku.hk

CP38

Selling Information Through Consulting

We consider a monopoly information holder selling information to a budget-constrained decision maker. The decision maker has a utility function that depends on his action and an uncertain state of the world. The seller and the buyer each observe a private signal regarding the state of the world. The seller's goal is to sell her private information to the buyer and extract maximum possible revenue, subject to the buyer's budget constraints. We consider three different settings, i.e., the seller and buyer signals can be independent, correlated, or follow a distribution accessed through a black-box sampling oracle. For each setting, we design information selling mechanisms which are both optimal and simple in the sense that they can be naturally interpreted, have succinct representations, and can be efficiently computed. Notably, all our mechanisms share the same format of acting as a consultant who recommends the best action to the buyer but uses different and carefully designed payment rules for different settings. Each of our optimal mechanisms can be easily computed by solving a single polynomial-size linear program. This significantly simplifies exponential-size LPs solved by the Ellipsoid method in the previous work, which computes the optimal mechanisms in the same setting but without budget limit. Such simplification is enabled by our new characterizations of the optimal mechanism in the budget-constrained setting.

Yiling Chen
Harvard University
yiling@seas.harvard.edu

Haifeng Xu
University of Virginia
hx4ad@virginia.edu

Shuran Zheng
Harvard University
Harvard University
shuran_zheng@seas.harvard.edu

CP39

Online Flow Computation on Unit-Vertex-Capacitated Networks

In many networking scenarios, long-lived flows can be rerouted to free up resources and accommodate new flows, but doing so comes at a cost in terms of disruption. An archetypical example is the transmission of live streams in a content delivery network: audio and video encoders (clients) generate live streams and connect to a server which rebroadcasts their stream to the rest of the network. Reconnecting a client to a different server mid-stream is very disruptive. We abstract these scenarios in the setting of a capacitated network where clients arrive one by one and request to send a unit of flow to a designated set of servers subject to edge/vertex capacity constraints. An online algorithm maintains a sequence of flows that route the clients present so far to the set of servers. The cost of a sequence of flows is defined as the net switching cost, i.e. total length of all augmenting paths used to transform each flow into its successor. We prove that for unit-vertex-capacitated networks, the algorithm that successively updates the flow using the shortest augmenting path from the new client to a free server incurs a total switching cost of $O(n \log^2 n)$, where n is the number of vertices in the network. This result is obtained by reducing to the online bipartite matching problem studied in prior work and applying their result. Finally, we identify a slightly more general class of networks for which essentially the same reduction idea can be applied to get the same bound.

Makis Arsenis, Robert Kleinberg
Cornell University
marsenis@gmail.com, rdk@cs.cornell.edu

CP39

Improved Parallel Cache-Oblivious Algorithms for Dynamic Programming

Emerging non-volatile main memory (NVRAM) technologies provide byte-addressability, low idle power, and improved memory-density, and are likely to be a key component in the future memory hierarchy. However, a critical challenge in achieving high performance is in accounting for the asymmetry that NVRAM writes can be significantly more expensive than NVRAM reads. In this paper, we consider a large class of cache-oblivious algorithms for dynamic programming (DP) and try to reduce the writes in the asymmetric setting while maintaining high parallelism. To achieve that, our key approach is to show the correspondence between these problems and an abstraction for their computation, which is referred to as the k-d grids. Then by showing lower bound and new algorithms for computing k-d grids, we show a list of improved cache-oblivious algorithms of many DP recurrences in the asymmetric setting, both sequentially and in parallel. Surprisingly, even without considering the read-write asymmetry (i.e., setting the write cost to be the same as the read cost in the algorithms), the new algorithms improve the existing cache complexity of many problems. We believe the reason is that the extra level of abstraction of k-d grids helps us to better understand the complexity and difficulties of these problems. We believe that the novelty of our framework is of interests and leads to many new questions for future work.

Yan Gu
University of California, Riverside
yan.gu@cs.cmu.edu

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

CP39

Fast Distributed Backup Placement in Sparse and Dense Networks

We consider the Backup Placement problem in networks in the CONGEST distributed setting. Given a network graph $G = (V, E)$, the goal of each vertex $v \in V$ is selecting a neighbor, such that the maximum number of vertices in V that select the same vertex is minimized. The backup placement problem was introduced by Halldorsson et al., who obtained an $O(\log n / \log \log n)$ approximation with randomized polylogarithmic time. In this paper we obtain significantly improved algorithms for various graph topologies. Specifically, we show that $O(1)$-approximation to optimal backup placement can be computed deterministically in $O(1)$ rounds in graphs that model wireless networks and more generally, in any graph with neighborhood independence bounded by a constant. At the other end, we consider sparse graphs, such as trees, forests, planar graphs and graphs of constant arboricity, and obtain a constant approximation to optimal backup placement in $O(\log n)$ deterministic rounds. Clearly, our constant-time algorithms for graphs with constant neighborhood independence are asymptotically optimal. Moreover, we show that our algorithms for sparse graphs are not far from optimal as well, by proving several lower bounds. Specifically, optimal backup placement of unoriented trees requires $\Omega(\log n)$ time, and approximate backup placement with a polylogarithmic approximation factor requires $\Omega(\sqrt{\log n/\log\log n})$ time.

Gal Oren
Ben-Gurion University of the Negev
Nuclear Research Center - Negev
orenw@post.bgu.ac.il

Leonid Barenboim
Open University of Israel
leonidb@openu.ac.il

CP39

Two-Way Trees: A Distributed Algorithm for Efficient Replica Search and Placement

We revisit a distributed caching protocol called random trees for hot spot relieving and low latency in large-scale networks. We propose a new version of random trees called *two-way (random) trees* that improve the logarithmic lookup path length $\Theta(\log_d N)$ to $\Theta(\sqrt{\log_d N})$ with high probability, where d is the degree and N is the size of the tree. Yet, all the optimal properties of random trees are preserved such as balanced workload and the minimum storage requirement without additional communication or computational cost. Two-way trees achieve these by separating trees for lookup and replication. A two-way tree is constructed in a fully distributed and maintenance-free manner, without *a-priori* known file popularity distributions. We present theoretical models to analyze the maximum workload of any server, and provide provable bounds comparable to the maximum load of a well-known balls into bins problem and a natural queueing model. The experimental results show that two-way trees reduce lookup path length by 60-70% compared to random trees, with minimum increase on server load and the number of replicas created.

Gahyun Park, Minseok Kwon, Ramprasad Tamilselvan
Rochester Institute of Technology
gxpics@rit.edu, jmk@cs.rit.edu, rt7516@rit.edu

Seungjoon Lee
Google
seungjoon@gmail.com

CP39

Eccentricity Heuristics Through Sublinear Analysis Lenses

The eccentricity of a node in a graph is its maximal shortest-path distance to any other node. Computing all eccentricities is a basic task in large-scale graph mining. Shun (KDD 2015) empirically studied two simple heuristics for this task: k-BFS-1, based on parallel BFS from a small sample of nodes, was shown to work well on a variety of graphs; k-BFS-2, a two-phase version, was shown to outperform state-of-the-art algorithms by up to orders of magnitude. This empirical success stands in apparent contrast to recent theoretical hardness results on approximating all eccentricities (Backurs et al., STOC 2018). This paper aims to formally explain the performance of these heuristics, by studying them through computational models designed for sublinear time or space algorithms. We use the proposed framework to derive improved variants, which retain their practicality while having better performance and formal guarantees. For k-BFS-1 we use property testing to match its performance with a more efficient variant. For k-BFS-2, we draw a connection to streaming Set Cover, and use it to suggest a variant which is guaranteed to compute almost all eccentricities exactly, if the graph satisfies a condition we call "small eccentric cover". The condition can be ascertained for all real-world graph used in Shun (KDD 2015) and in our experiments. Experimental results on real-world graphs demonstrate the validity of our analysis and the empirical advantage of the proposed methods.

Tal Wagner
MIT
MIT
tal.wagner@gmail.com

CP40

Fully Dynamic Matching: Beating 2-Approximation in Δ^ϵ Update Time

In fully dynamic graphs, we know how to maintain a 2-approximation of maximum matching extremely fast, that is, in polylogarithmic update time or better. In a sharp contrast and despite extensive studies, all known algorithms that maintain a $2-\Omega(1)$ approximate matching are much slower. Understanding this gap and, in particular, determining the best possible update time for algorithms providing a better-than-2 approximate matching is a major open question. In this paper, we show that for any constant $\epsilon > 0$, there is a randomized algorithm that with high probability maintains a $2-\Omega(1)$ approximate maximum matching of a fully-dynamic general graph in worst-case update time $O(\Delta^\epsilon + \text{polylog}\, n)$, where Δ is the maximum degree. Previously, the fastest fully dynamic matching algorithm providing a better-than-2 approximation had $O(m^{1/4})$ update-time [Bernstein and Stein, SODA 2016]. A faster algorithm with update-time $O(n^\epsilon)$ was known, but worked only for maintaining the size (and not the edges) of the matching in bipartite graphs [Bhattacharya, Henzinger, and Nanongkai, STOC 2016].

Soheil Behnezhad
University of Maryland
soheil@cs.umd.edu

Jakub Lacki
Google Research
jlacki@google.com

Vahab Mirrokni
Google Inc.

mirrokni@google.com

CP40

An Improved Algorithm for Incremental Cycle Detection and Topological Ordering in Sparse Graphs

We consider the problem of incremental cycle detection and topological ordering in a directed graph $G = (V, E)$ with $|V| = n$ nodes. In this setting, initially the edge-set E of the graph is empty. Subsequently, at each time-step an edge gets inserted into G. After every edge-insertion, we have to report if the current graph contains a cycle, and as long as the graph remains acyclic, we have to maintain a topological ordering of the node-set V. Let m be the total number of edges that get inserted into G. We present a randomized algorithm for this problem with $\tilde{O}(m^{4/3})$ total expected update time. Our result happens to be the first one that breaks the $\Omega(n\sqrt{m})$ lower bound of [**?**] on the total update time of any *local* algorithm for a nontrivial range of sparsity. Specifically, the total update time of our algorithm is $o(n\sqrt{m})$ whenever $m = o(n^{6/5})$. From a technical perspective, we obtain our result by combining the algorithm of [**?**] with the *balanced search* framework of [**?**].

Sayan Bhattacharya
University of Warwick
Coventry
S.Bhattacharya@warwick.ac.uk

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

CP40

Fully-Dynamic All-Pairs Shortest Paths: Improved Worst-Case Time and Space Bounds

Given a directed weighted graph $G = (V, E)$ undergoing vertex insertions *and* deletions, the All-Pairs Shortest Paths (APSP) problem asks to maintain a data structure that processes updates efficiently and returns after each update the distance matrix to the current version of G. In two breakthrough results, Italiano and Demetrescu [STOC '03] presented an algorithm that requires $\tilde{O}(n^2)$ *amortized* update time, and Thorup showed in [STOC '05] that *worst-case* update time $\tilde{O}(n^{2+3/4})$ can be achieved. In this article, we make substantial progress on the problem. We present the following new results:

- the first deterministic data structure that breaks the $\tilde{O}(n^{2+3/4})$ worst-case update time bound by Thorup which has been standing for almost 15 years. We improve the worst-case update time to $\tilde{O}(n^{2+5/7}) = \tilde{O}(n^{2.71..})$ for weighted and to $\tilde{O}(n^{2+3/5}) = \tilde{O}(n^{2.6})$ for unweighted graphs.
- a simple deterministic algorithm with $O(n^{2+3/4})$ worst-case update time ($O(n^{2+2/3})$ for unweighted graphs), and a simple Las-Vegas algorithm with worst-case update time $\tilde{O}(n^{2+2/3})$ ($\tilde{O}(n^{2+1/2})$ for unweighted graphs) that works against a non-oblivious adversary. Both data structures require space $\tilde{O}(n^2)$. These are the first exact dynamic algorithms with truly-subcubic update time *and* space usage.

Maximilian Probst Gutenberg
University of Copenhagen
probst@di.ku.dk

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen
koolooz@di.ku.dk

CP40

Deterministic Algorithms for Decremental Approximate Shortest Paths: Faster and Simpler

In the decremental $(1 + \epsilon)$-approximate Single-Source Shortest Path (SSSP) problem, we are given an undirected, unweighted graph $G = (V, E)$ with $n = |V|, m = |E|$, undergoing edge deletions, and a distinguished source $s \in V$, and we are asked to process edge deletions efficiently and answer queries for distance estimates $\widetilde{\mathbf{dist}}_G(s, v)$ for each $v \in V$, at any stage, such that $\mathbf{dist}_G(s, v) \leq \widetilde{\mathbf{dist}}_G(s, v) \leq (1 + \epsilon)\mathbf{dist}_G(s, v)$. In the decremental $(1 + \epsilon)$-approximate All-Pairs Shortest Path (APSP) problem, we are asked to answer queries for distance estimates $\widetilde{\mathbf{dist}}_G(u, v)$ for every $u, v \in V$. We present a new *deterministic* algorithm for the decremental $(1 + \epsilon)$-approximate SSSP problem that takes total update time $O(mn^{0.5+o(1)})$. Our algorithm improves on the currently best algorithm for dense graphs by Chechik and Bernstein [STOC 2016] with total update time $\tilde{O}(n^2)$ and the best existing algorithm for sparse graphs with running time $\tilde{O}(n^{1.25}\sqrt{m})$ [SODA 2017] whenever $m = O(n^{1.5-o(1)})$. As a by-product, we also obtain a new simple deterministic algorithm for the decremental $(1 + \epsilon)$-approximate APSP problem with near-optimal total running time $\tilde{O}(mn/\epsilon)$ matching the time complexity of the sophisticated but rather involved algorithm by Henzinger, Forster and Nanongkai [FOCS 2013].

Maximilian Probst Gutenberg
University of Copenhagen
probst@di.ku.dk

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen
koolooz@di.ku.dk

CP40

Decremental SSSP in Weighted Digraphs: Faster and Against an Adaptive Adversary

Given a dynamic digraph $G = (V, E)$ undergoing edge deletions and given $s \in V$ and constant $\epsilon > 0$, we consider the problem of maintaining $(1 + \epsilon)$-approximate shortest path distances from s to all vertices in G over the sequence of deletions. Recently, Henzinger et al. (STOC'14, ICALP'15) presented an algorithm with total update time $O(mn^{0.9+o(1)} \log W)$ breaking the longstanding $O(mn)$ deterministic algorithm by Even and Shiloach (J. ACM'81). Here, we present the following $(1 + \epsilon)$-approximate data structures:

1. a Las Vegas data structure that works against an adaptive adversary with total expected update time $\tilde{O}(m^{2/3}n^{4/3})$ for unweighted graphs and $\tilde{O}(m^{3/4}n^{5/4} \log W)$ for weighted graphs,
2. a Las Vegas data structure with total expected update time $\tilde{O}(\sqrt{m}n^{3/2})$ for unweighted graphs and $\tilde{O}(m^{2/3}n^{4/3} \log W)$ for weighted graphs,
3. a Monte Carlo data structure with total expected update time $\tilde{O}((mn)^{7/8} \log W) = \tilde{O}(mn^{3/4} \log W)$.

Combined, our update times are faster than those of Henzinger et al. for all graph densities. Furthermore, our first data structure is the first to break the $O(mn)$ bound while still working against an adaptive adversary.

Maximilian Probst Gutenberg
University of Copenhagen
probst@di.ku.dk

Christian Wulff-Nilsen
Department of Computer Science
University of Copenhagen
koolooz@di.ku.dk

CP41

Hierarchical Shape Construction and Complexity for Slidable Polyominos under Uniform External Forces

Advances in technology have given us the ability to create and manipulate robots for numerous applications at the molecular scale. At this size, fabrication tool limitations motivate the use of simple robots. The individual control of these simple objects can be infeasible. We investigate a model of robot motion planning, based on global external signals, known as the tilt model. Given a board and initial placement of polyominoes, 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. We propose a new hierarchy of shapes and design a single configuration that is strongly universal for any $w \times h$ bounded shape within this hierarchy (it can be reconfigured to any $w \times h$ bounded shape in the hierarchy). This class of shapes constitutes the most general set of buildable shapes in the literature, with most previous work consisting of just the first-level of our hierarchy. We accompany this result with a $O(n^4 \log n)$-time algorithm for deciding if a given hole-free shape is a member of the hierarchy. For our second result, we resolve a long-standing open problem within the field: deciding if a given position may be covered by a tile for a given initial board configuration is PSPACE-complete, even when all movable pieces are 1×1 tiles with no glues. We achieve this with a reduction from Non-deterministic Constraint Logic for a one-player unbounded game.

Tim Wylie
University of Texas - Rio Grande Valley
timothy.wylie@utrgv.edu

Jose Balanza-Martinez, David Caballero, Angel Cantu, Mauricio Flores, Timothy Gomez, Austin Luchsinger, Rene Reyes, Robert Schweller
University of Texas Rio Grande Valley
jose.balanzamartinez01@utrgv.edu, david.caballero01@utrgv.edu, angel.cantu01@utrgv.edu, mauricio.flores01@utrgv.edu, timothy.gomez01@utrgv.edu, austin.luchsinger01@utrgv.edu, rene.reyes01@utrgv.edu, robert.schweller@utrgv.edu

CP41

The Impacts of Dimensionality, Diffusion, and Directedness on Intrinsic Universality in the Abstract Tile Assembly Model

We present results related to mathematical models of self-assembling tiles and the impacts 3 diverse properties have on their dynamics. We expand upon prior results showing that (1) the abstract Tile Assembly Model (aTAM) is intrinsically universal (IU) [FOCS 2012], and (2) the class of directed aTAM systems is not IU [FOCS 2016]. IU for a model (or class of systems in a model) means there's a universal tile set that can simulate an arbitrary system within it. Furthermore, the simulation must not only produce the same resultant structures, it must also maintain the full dynamics of the systems being simulated modulo only a scale factor. While the FOCS 2012 result showed the 2D aTAM is IU, here we show this is also the case for the 3D version. Conversely, the FOCS 2016 result showed the class of aTAM systems which are directed (a.k.a. deterministic, or confluent) is not IU, implying that nondeterminism is fundamentally required for such simulations. Here, however, we show that in 3D the class of directed aTAM systems is IU, i.e. there is a universal directed simulator for them. We then consider the influence of more rigid notions of dimensionality. We introduce the Planar aTAM, where tiles are not only restricted to binding in the plane, but also to traveling in the plane, and prove it is not IU, and that the class of directed systems within it also is not IU. Finally, analogous to the Planar aTAM, we introduce the Spatial aTAM, its 3D counterpart, and prove that it is IU.

Daniel Hader, Aaron Koch, Matthew Patitz, Michael Sharp
University of Arkansas
dhader@uark.edu, aekoch@uark.edu, patitz@uark.edu, mrs018@uark.edu

CP41

Computing Minimal Persistent Cycles: Polynomial and Hard Cases

Persistent cycles, especially the minimal ones, are useful geometric features functioning as augmentations for the intervals in the purely topological persistence diagrams (also termed as barcodes). In our earlier work, we showed that computing minimal 1-dimensional persistent cycles (persistent 1-cycles) for finite intervals is NP-hard while the same for infinite intervals is polynomially tractable. In this paper, we address this problem for general dimensions with Z_2 coefficients. In addition to proving that it is NP-hard to compute minimal persistent d-cycles (d>1) for both types of intervals given arbitrary simplicial complexes, we identify two interesting cases which are polynomially tractable. These two cases assume the complex to be a certain generalization of manifolds which we term as weak pseudomanifolds. For finite intervals from the d-th persistence diagram of a weak (d+1)-pseudomanifold, we utilize the fact that persistent cycles of such intervals are null-homologous and reduce the problem to a minimal cut problem. Since the same problem for infinite intervals is NP-hard, we further assume the weak (d+1)-pseudomanifold to be embedded in $\mathbb{R}^{d+1}$ so that the complex has a natural dual graph structure and the problem reduces to a minimal cut problem. Experiments with both algorithms on scientific data indicate that the minimal persistent cycles capture various significant features of the data.

Tamal K. Dey
The Ohio State University
Columbus, Ohio 43210
dey.8@osu.edu

Tao Hou, Sayan Mandal
Department of Computer Science and Engineering
The Ohio State University

hou.332@osu.edu, mandal.25@osu.edu

CP41

Even Maps, the Colin de Verdire Number and Representations of Graphs

Van der Holst and Pendavingh introduced a graph parameter σ, which coincides with the more famous Colin de Verdière graph parameter μ for small values. However, the definition of σ is much more geometric/topological directly reflecting embeddability properties of the graph. They proved $\mu(G) \leq \sigma(G) + 2$ and conjectured $\mu(G) \leq \sigma(G)$ for any graph G. We confirm this conjecture. As far as we know, this is the first topological upper bound on $\mu(G)$ which is, in general, tight. Equality between μ and σ does not hold in general as van der Holst and Pendavingh showed that there is a graph G with $\mu(G) \leq 18$ and $\sigma(G) \geq 20$. We show that the gap appears on much smaller values, namely, we exhibit a graph H for which $\mu(H) \leq 7$ and $\sigma(H) \geq 8$. We also prove that, in general, the gap can be large: The incidence graphs H_q of finite projective planes of order q satisfy $\mu(H_q) \in O(q^{3/2})$ and $\sigma(H_q) \geq q^2$.

Vojtech Kalua
Institut für Mathematik
Universität Innsbruck
vojtech.kaluza@uibk.ac.at

Martin Tancer
Charles University in Prague
tancer@kam.mff.cuni.cz

CP41

On the Cover of the Rolling Stone

We construct a convex polytope of unit diameter that can be placed on a horizontal surface on one of its faces, so that it will repeatedly roll over from one face to another until it comes to rest on another face, and covering an arbitrarily large distance in this process: that is, the horizontal distance between the footprints of the start and final faces can be larger than any given threshold. The movement is solely caused by the center of mass of the polytope falling outside of the face the polytope is standing on before each move, while the speed of the motion is irrelevant. That is, if the polytope is manually stopped after each tumble, the motion will resume when released (unless it stands on the final stable face). According to the law of physics, the vertical distance between the center of mass of the polytope and the horizontal surface continuously decreases throughout the entire motion. Moreover, our polytope can be realized so that (i) it has a unique stable facet, and (ii) it is an arbitrary close approximation of a unit ball. As such, this construction gives a positive answer to a question raised by Conway (1969). The arbitrarily large rolling distance property investigated here for the first time raises intriguing questions and opens new avenues for future research.

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

Csaba Toth
Department of Mathematics
California State University Northridge
cdtoth@acm.org

CP42

A Little Charity Guarantees Almost Envy-Freeness

Fair division of indivisible goods is a very well-studied problem. The goal of this problem is to distribute m goods to n agents in a “fair’ manner, where every agent has a valuation for each subset of goods. We assume general valuations. Envy-freeness is the most extensively studied notion of fairness. However, envy-free allocations do not always exist when goods are indivisible. The notion of fairness we consider here is “envy-freeness up to any good’ (EFX) where no agent envies another agent after the removal of any single good from the other agent’s bundle. It is not known if such an allocation always exists even when $n = 3$. We show there is always a partition of the set of goods into $n + 1$ subsets $(X_1, \ldots, X_n, P)$ where for $i \in [n]$, X_i is the bundle allocated to agent i and the set P is unallocated (or donated to charity) such that we have:

- envy-freeness up to any good,
- no agent values P higher than her own bundle, and
- fewer than n goods go to charity, i.e., $|P| < n$ (typically $m \gg n$).

Our proof is constructive. Moreover, a minor variant of our algorithm also shows the existence of an allocation which is 4/7 *Groupwise* MMS (GMMS): this is a notion of fairness stronger than MMS and EFX. This improves upon the current best bound of 1/2 known for an approximate GMMS allocation.

Bhaskar R. Chaudhury
Max Planck Institute for Informatics
braycha@mpi-inf.mpg.de

Tellikepalli Kavitha
Tata Institute of Fundamental Research
kavitha@tifr.res.in

Kurt Mehlhorn
MPI for Informatics
Saarbruecken, Germany
mehlhorn@mpi-inf.mpg.de

Alkmini Sgouritsa
University of Liverpool
alkmini@liv.ac.uk,

CP42

The Complexity of Contracts

We initiate the study of computing (near-)optimal contracts in succinctly representable principal-agent settings. Here optimality means maximizing the principal’s expected payoff over all incentive-compatible contracts—known in economics as “second-best” solutions. We also study a natural relaxation to approximately incentive-compatible contracts. We focus on principal-agent settings with succinctly described (and exponentially large) outcome spaces. We show that the computational complexity of computing a near-optimal contract depends fundamentally on the number of agent actions. For settings with a constant number of actions, we present a fully polynomial-time approximation scheme (FPTAS) for the separation oracle of the dual of the problem of minimizing the principal’s payment to the agent, and use this subroutine to efficiently compute a δ-incentive-compatible (δ-IC) contract whose expected payoff matches or surpasses that of the optimal IC contract. With an arbitrary number of actions, we prove that the problem is hard to approximate within any constant c. This inapproximability result holds even for δ-IC contracts where δ is a sufficiently slow-growing function of c. On the positive side, we show that simple linear δ-IC contracts with constant δ are sufficient to achieve a constant-factor approximation of the "first-best" (full-welfare-extracting) solution, and that such a contract can be computed in polynomial time.

Paul Dütting
London School of Economics
p.d.duetting@lse.ac.uk

Tim Roughgarden
Columbia University
tr@cs.columbia.edu

Inbal Talgam-Cohen
Technion
italgam@cs.technion.ac.il

CP42

Approximating Nash Social Welfare under Submodular Valuations Through (Un)Matchings

We study the problem of approximating maximum Nash social welfare (NSW) when allocating m indivisible items among n asymmetric agents with submodular valuations. The NSW is a well-established notion of fairness and efficiency, defined as a weighted geometric mean of agents' valuations. For special cases with symmetric agents and additive(-like) valuations, approximation algorithms have been designed using approaches customized for these specific settings, and no known algorithm gives better than $\Omega(m)$ factor either for asymmetric agents with additive valuations or symmetric agents beyond additive(-like) valuations. We present two approximation algorithms for asymmetric agents with additive and submodular valuations. Both algorithms involve non-trivial modifications of a greedy repeated matchings approach, where allocations of high valued items are done separately by un-matching certain items and re-matching them in different ways. We show that the algorithms achieve approximation factors of $O(n)$ and $O(n \log n)$ for additive and submodular valuations respectively. For additive valuations, our algorithm's output also has the fairness property of envy-free up to one item. We also show that the NSW problem under submodular valuations is strictly harder than all known settings with a e/(e-1) factor of hardness of approximation, even for constantly many agents. For this case, we provide a different approximation algorithm that achieves a matching factor of e/(e-1).

Rucha Kulkarni, Jugal Garg
University of Illinois at Urbana-Champaign
ruchark2@illinois.edu, jugal@illinois.edu

Pooja Kulkarni
University of Illinois at Urbana Champaign
poojark2@illinois.edu

CP42

Complexity and Parametric Computation of Equilibria in Atomic Splittable Congestion Games Via Weighted Block Laplacians

We show that computing an equilibrium in atomic splittable congestion games with player-specific affine cost functions $l_{e,i}(x) = a_{e,i}x + b_{e,i}$ is PPAD-complete. To prove that the problem is contained in PPAD, we develop a homotopy method that traces an equilibrium for varying flow demands of the players. A key technique for this method is to describe the evolution of the equilibrium locally by a novel block Laplacian matrix. Using the properties of this matrix give rise to a path following formulation for computing an equilibrium where states correspond to supports that are feasible for certain demands. For the PPAD-hardness, we reduce from computing an approximate equilibrium of a bimatrix win-lose game. On our way, we obtain new results regarding the multiplicity of equilibria. When the coefficients $a_{e,i}$ are in general position, every game has a finite set of equilibria while without this assumption there may be a continuum of equilibria. When the additive constants $b_{e,i}$ are in general position, games have an odd number of equilibria except for a nullset of demand values. As another byproduct of our analysis, we obtain an algorithm that computes a continuum of equilibria parametrized by the players' flow demand. For player-specific costs, the algorithm runs in polynomial space. For games with player-independent costs, we obtain an algorithm computing all equilibria as a function of the flow demand that runs in time polynomial in the output.

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

CP42

On the Tractability of Public Persuasion with No Externalities

Persuasion studies how a principal can influence agents' decisions via strategic information revelation — often described as a *signaling scheme* — in order to yield the most desirable equilibrium outcome. A basic question that has attracted much recent attention is how to compute the optimal public signaling scheme, a.k.a., *public persuasion*, which is motivated by various applications including auction design, routing, voting, marketing, queuing, etc. Unfortunately, most algorithmic studies in this space exhibit quite negative results and are rifle with computational intractability. Given such background, this paper seeks to understand when public persuasion is tractable and how tractable it can be. We focus on a fundamental multi-agent persuasion model introduced by Arieli and Babichenko: many agents, *no* inter-agent externalities and binary agent actions, and identify well-motivated circumstances under which efficient algorithms are possible. En route, we also develop new algorithmic techniques and demonstrate that they can be applicable to other public persuasion problems or even beyond.

Haifeng Xu
University of Virginia
hx4ad@virginia.edu

CP43

Cilkmem: Algorithms for Analyzing the Memory High-Water Mark of Fork-Join Parallel Programs

Software engineers designing recursive fork-join programs destined to run on massively parallel computing systems must be cognizant of how their program's memory requirements scale in a many-processor execution. Although tools exist for measuring memory usage during one particular execution of a parallel program, such tools cannot bound the worst-case memory usage over all possible parallel executions. This paper introduces Cilkmem, a tool that analyzes the execution of a deterministic Cilk program to determine its p-processor memory high-water mark (MHWM), which is the worst-case memory usage of the program over *all possible* p-processor executions. Cilkmem employs two new algorithms for computing the p-processor MHWM. The first algorithm calculates the exact p-processor MHWM in $O(T_1 \cdot p)$ time, where T_1 is the total work of the program. The second algorithm solves, in $O(T_1)$ time, the approximate threshold problem, which asks, for a given memory threshold M, whether the p-processor MHWM exceeds $M/2$ or whether it is guaranteed to be less than M. Across ten application benchmarks from the Cilkbench suite, the exact algorithm incurs a geometric-mean multiplicative overhead of 1.54 for $p = 128$, whereas the approximation-threshold algorithm incurs an overhead of 1.36 independent of p. In addition, we use Cilkmem to diagnose a previously unknown issue in a large image-alignment program contributing to unexpectedly high memory usage.

Tim Kaler
Massachusetts Institute of Technology (CSAIL)
tfk@mit.edu

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

Tao Schardl, Daniele Vettorel
Massachusetts Institute of Technology (CSAIL)
neboat@mit.edu, vettorel@mit.edu

CP43

Scheduling I/O Latency-Hiding Futures in Task-Parallel Platforms

Task parallelism research has traditionally focused on optimizing computation-intensive applications. Due to the proliferation of commodity parallel processors, there has been recent interest in supporting interactive applications. Such interactive applications frequently rely on I/O operations that require few processing cycles but may incur significant latency to complete. In order to increase performance, when a particular thread of control is blocked on an I/O operation, ideally we would like to hide this latency by using the processing resources to do other ready work instead of blocking or spin waiting on this I/O. There has been limited prior work on hiding this latency and only one result that provides a theoretical bound for interactive applications that use I/O operations. In this work, we propose a task parallel platform that supports I/O operations using the futures abstraction and a corresponding scheduler that schedules the I/O operations while hiding their latency. We provide a theoretical analysis of our scheduling algorithm that shows our algorithm provides better execution time guarantees than prior work. We also implemented the algorithm in a practically efficient prototype library that runs on top of the Cilk-F runtime, a runtime system that supports futures within the context of the Cilk Plus language, and performed experiments that demonstrate the efficiency of our implementation.

Kyle Singer, Kunal Agrawal, I-Ting Angelina Lee
Washington University in St. Louis
kdsinger@wustl.edu, kunal@wustl.edu, angelee@wustl.edu

CP44

A Tale of Santa Claus, Hypergraphs and Matroids

A well-known problem in scheduling and approximation algorithms is the Santa Claus problem. Suppose that Santa Claus has a set of gifts, and he wants to distribute them among a set of children so that the least happy child is made as happy as possible. Here, the value that a child i has for a present j is of the form $p_{ij} \in \{0, p_j\}$. A polynomial time algorithm by Annamalai et al. gives a 12.33-approximation and is based on a modification of Haxell's hypergraph matching argument. In this paper, we introduce a *matroid* version of the Santa Claus problem. Our algorithm is also based on Haxell's augmenting tree, but with the introduction of the matroid structure we solve a more general problem with cleaner methods. Our result can then be used as a blackbox to obtain a $(4+\varepsilon)$-approximation for Santa Claus. This factor also compares against a natural, compact LP for Santa Claus.

Sami Davies, Thomas Rothvoss, Yihao Zhang
University of Washington
daviess@uw.edu, rothvoss@uw.edu, yihaoz93@uw.edu

CP44

Weighted Completion Time Minimization for Unrelated Machines Via Iterative Fair Contention Resolution

We give a 1.488-approximation for the classic scheduling problem of minimizing total weighted completion time on unrelated machines. This is a considerable improvement on the recent breakthrough of $(1.5 - 10^{-7})$-approximation (STOC 2016, Bansal-Srinivasan-Svensson) and the follow-up result of $(1.5 - 1/6000)$-approximation (FOCS 2017, Li). Bansal et al. introduced a novel rounding scheme yielding strong negative correlations for the first time and applied it to the scheduling problem to obtain their breakthrough, which resolved the open problem if one can beat out the long-standing 1.5-approximation barrier based on independent rounding. Our key technical contribution is in achieving significantly stronger negative correlations via *iterative fair contention resolution*, which is of independent interest. Previously, Bansal et al. obtained strong negative correlations via a variant of pipage type rounding and Li used it as a black box.

Sungjin Im, Maryam Shadloo
University of California, Merced
sim3@ucmerced.edu, mshadloo@ucmerced.edu

CP44

Parallel Machine Scheduling to Minimize Energy Consumption

Given n jobs with release dates, deadlines and processing times we consider the problem of scheduling them on m parallel machines so as to minimize the total energy consumed. Machines can enter a sleep state and they consume no energy in this state. Each machine requires Q units of energy to awaken from the sleep state and in its active state the machine can process jobs and consumes a unit of energy per unit time. We allow for preemption and migration of jobs and provide the first constant approximation algorithm for this problem.

Nikhil Kumar

Indian Institute of Technology (IIT) Delhi
nikhil@cse.iitd.ac.in

Antonios Antoniadis
Max-Planck-Institut für Informatik
aantonia@mpi-inf.mpg.de

Naveen Garg
IIT Delhi
naveen@cse.iitd.ac.in

Gunjan Kumar
Tata Institute of Fundamental Research, Mumbai
gunjan.kumar@tifr.res.in

CP44

How to Aggregate Top-Lists: Approximation Algorithms Via Scores and Average Ranks

A top-list is a possibly incomplete ranking of elements: only a subset of the elements are ranked, with all unranked elements tied for last. Top-list aggregation, a generalization of the well-known rank aggregation problem, takes as input a collection of top-lists and aggregates them into a single complete ranking, aiming to minimize the number of upsets (pairs ranked in opposite order in the input and in the output). In this paper, we give simple approximation algorithms for top-list aggregation.

- We generalize the footrule algorithm for rank aggregation.
- Using inspiration from approval voting, we define the score of an element as the frequency with which it is ranked, i.e. appears in an input top-list. We reinterpret Ailon's RepeatChoice algorithm for top-list aggregation using the score of an element and its average rank given that it is ranked.
- Using average ranks, we generalize and analyze Borda's algorithm for rank aggregation.
- We design a simple 2-phase variant of the Generalized Borda's algorithm, roughly sorting by scores and breaking ties by average ranks.
- We then design another 2-phase variant in which in order to break ties we use, as a black box, the Mathieu-Schudy PTAS for rank aggregation, yielding a PTAS for top-list aggregation.
- Finally, we discuss the special case in which all input lists have constant length.

Simon Mauras
Université de Paris, IRIF, CNRS, F-75013 Paris
simon.mauras@irif.fr

Claire Mathieu
CNRS, IRIF
clairemmathieu@gmail.com

CP44

Hierarchy-Based Algorithms for Minimizing Makespan under Precedence and Communication Constraints

We consider the classic problem of scheduling jobs with precedence constraints on a set of identical machines to minimize the makespan objective function. Understanding the exact approximability of the problem when the number of machines is a constant is a well-known question in scheduling theory. Indeed, an outstanding open problem from the classic book of Garey and Johnson asks whether this problem is NP-hard even in the case of 3 machines and unit-length jobs. In a recent breakthrough, Levey and Rothvoss gave a (1+epsilon)- approximation algorithm, which runs in nearly quasi-polynomial time, for the case when jobs have unit lengths. However, a substantially more difficult case where jobs have arbitrary processing lengths has remained open. We make progress on this more general problem. We show that there exists a (1+epsilon)-approximation algorithm (with similar running time as that of Levey and Rothvoss) for the non-migratory setting: when every job has to be scheduled entirely on a single machine, but within a machine the job need not be scheduled during consecutive time steps. Further, we also show that our algorithmic framework generalizes to another classic scenario where, along with the precedence constraints, the jobs also have communication delay constraints. Both of these fundamental problems are highly relevant to the practice of datacenter scheduling.

Janardhan Kulkarni
Microsoft Research, Redmond
jakul@microsoft.com

Shi Li
University at Buffalo
shil@buffalo.edu

Jakub Tarnawski
MSR, Redmond
jakub.tarnawski@epfl.ch

Minwei Ye
University at Buffalo
minweiye@buffalo.edu

CP45

Computing Circle Packing Representations of Planar Graphs

The Circle Packing Theorem states that every planar graph can be represented as the tangency graph of a family of internally-disjoint circles. A well-known generalization is the Primal-Dual Circle Packing Theorem for 3-connected planar graphs. The existence of these representations has widespread applications in theoretical computer science and discrete mathematics; however, the algorithmic aspect has received relatively little attention. In this work, we present an algorithm based on convex optimization for computing a primal-dual circle packing representation of maximal planar graphs, i.e. triangulations. This in turn gives an algorithm for computing a circle packing representation of any planar graph. Both take $\widetilde{O}(n\log(R/\varepsilon))$ expected run-time to produce a solution that is ε close to a true representation, where R is the ratio between the maximum and minimum circle radius in the true representation.

Sally Dong, Yin Tat Lee
University of Washington
sallyqd@uw.edu, yintat@uw.edu

Kent Quanrud
University of Illinois Urbana-Champaign

quanrud2@illinois.edu

CP45

Atomic Embeddability, Clustered Planarity, and Thickenability

We study the atomic embeddability testing problem, which is a common generalization of clustered planarity (c-planarity, for short) and thickenability testing, and present a polynomial time algorithm for this problem, thereby giving the first polynomial time algorithm for c-planarity. C-planarity was introduced in 1995 by Feng, Cohen, and Eades as a variant of graph planarity, in which the vertex set of the input graph is endowed with a hierarchical clustering and we seek an embedding (crossing free drawing) of the graph in the plane that respects the clustering in a certain natural sense. Until now, it has been an open problem whether c-planarity can be tested efficiently, despite relentless efforts. The thickenability problem for simplicial complexes emerged in the topology of manifolds in the 1960s. A 2-dimensional simplicial complex is thickenable if it embeds in some orientable 3-dimensional manifold. Recently, Carmesin announced that thickenability can be tested in polynomial time. Our algorithm for atomic embeddability combines ideas from Carmesin's work with algorithmic tools previously developed for weak embeddability testing. We express our results purely in terms of graphs on surfaces, and rely on the machinery of topological graph theory. Finally we give a polynomial-time reduction from c-planarity to thickenability and show that a slight generalization of atomic embeddability to the setting in which clusters are toroidal graphs is NP-complete.

Radoslav Fulek
Institute of Science and Technology Austria
radoslav.fulek@gmail.com

CP45

Nearly Optimal Planar k Nearest Neighbors Queries under General Distance Functions

We study the k nearest neighbors problem in the plane for general, convex, pairwise disjoint sites of constant description complexity such as line segments, disks, and quadrilaterals and with respect to a general family of distance functions including the L_p-norms and additively weighted Euclidean distances. For point sites in the Euclidean metric, after four decades of effort, an optimal data structure has recently been developed with $O(n)$ space, $O(\log n + k)$ query time, and $O(n \log n)$ preprocessing time. We develop a static data structure for the general setting with nearly optimal $O(n \log \log n)$ space, the optimal $O(\log n + k)$ query time, and expected $O(n\ polylog\ n)$ preprocessing time. Our dynamic version (that allows insertions and deletions of sites) requires $O(n \log n)$ space while achieving $O(\log^2 n + k)$ query time and $O(polylog\ n)$ update time, and thus improves many applications. To obtain these improvements, we devise shallow cuttings of linear size for general distance functions. Shallow cuttings are a key technique to deal with the k nearest neighbors problem for point sites in the Euclidean metric. Although shallow cuttings have been generalized to general distance functions, existing ones either could not be applied to the k nearest neighbors problem or have an extra polylog-factor in size. Our innovation is a new random sampling technique for the analysis of geometric structures, and we believe it is of independent interest.

Chih-Hung Liu
Department of Computer Science
ETH Zürich
chih-hung.liu@inf.ethz.ch

CP45

Connectivity of Triangulation Flip Graphs in the Plane (Part I: Edge Flips)

In a straight-line embedded triangulation of a point set P in the plane, removing an inner edge and—provided the resulting quadrilateral is convex—adding the other diagonal is called an *edge flip*. The *(edge) flip graph* has all triangulations as vertices, and a pair of triangulations is adjacent if they can be obtained from each other by an edge flip. The goal of this paper is to contribute to a better understanding of the flip graph, with an emphasis on its connectivity. For sets in general position, it is known that every triangulation allows at least $\lceil n2 - 2\rceil$ edge flips (a tight bound) which gives the minimum degree of any flip graph for n points. We show that for every point set P in general position, the flip graph is at least $\lceil n2 - 2\rceil$-vertex connected. Somewhat more strongly, we show that the vertex connectivity equals the minimum degree occurring in the flip graph, i.e. the minimum number of flippable edges in any triangulation of P, provided P is large enough. Furthermore, we show that the flip graph can be covered by 1-skeletons of polytopes of dimension $\lceil n2 - 2\rceil$ (products of associahedra). A corresponding result ($(n-3)$-vertex connectedness) can be shown for the bistellar flip graph of partial triangulations, i.e., the set of all triangulations of subsets of P which contain all extreme points of P. This will be treated separately in a second part.

Uli Wagner
IST Austria
uli@ist.ac.at

Emo Welzl
ETH Zürich
emo@inf.ethz.ch

CP45

The Stable Set Problem in Graphs with Bounded Genus and Bounded Odd Cycle Packing Number

Consider the family of graphs without k node-disjoint odd cycles, where k is a constant. Determining the complexity of the stable set problem for such graphs G is a long-standing problem. We give a polynomial-time algorithm for the case that G can be further embedded in a (possibly non-orientable) surface of bounded genus. Moreover, we obtain polynomial-size extended formulations for the respective stable set polytopes. To this end, we show that 2-sided odd cycles satisfy the Erdős-Pósa property in graphs embedded in a fixed surface. This extends the fact that odd cycles satisfy the Erdős-Pósa property in graphs embedded in a fixed orientable surface [Kawarabayashi and Nakamoto, The Erdős-Pósa property for vertex- and edge-disjoint oddcycles in graphs on orientable surfaces, Discrete Math., 307(6):764768, 2007]. Eventually, our findings allow us to reduce the original problem to the problem of finding a minimum-cost non-negative integer circulation of a certain homology class, which turns out to be efficiently solvable in our case.

Stefan Weltge
ETH Zurich
weltge@tum.de

Michele Conforti
Università degli Studi di Padova
conforti@math.unipd.it

Samuel Fiorini, Tony Huynh
Université Libre de Bruxelles
sfiorini@ulb.ac.be, tony.bourbaki@gmail.com

Gwenaël Joret
Département d'Informatique, Université Libre de Bruxelles
Brussels, Belgium
gjoret@ulb.ac.be

CP46

Sublinear Time Approximation of the Cost of a Metric k-Nearest Neighbor Graph

Let (X, d) be an n-point metric space. We assume that (X, d) is given in the distance oracle model. A k-nearest neighbor graph for (X, d) is a directed graph $G = (V, E)$ that has an edge to each of v's k nearest neighbors. We use $cost(G)$ to denote the sum of edge weights of G. We study the problem of approximating $cost(G)$ in sublinear time, when we are given oracle access to the metric space (X, d) defining G. Our goal is to develop an algorithm that solves this problem faster than the time required to compute G. We first present an algorithm that in $\tilde{O}(n^2/k)$ time with probability at least $\frac{2}{3}$ approximates $cost(G)$ to within a factor of $1 + \epsilon$. Next, we present a more elaborate sublinear algorithm that in time $\tilde{O}(\min\{nk^{3/2}, n^2/k\})$ computes an estimate $cost'(G)$ of $cost(G)$ that with probability at least $\frac{2}{3}$ satisfies $|cost(G) - cost'(G))| \leq \epsilon \cdot (cost(G) + MST(X))$, where $MST(X)$ is the cost of the MST of (X, d). Further, we complement these results with near matching lower bounds. We show that any algorithm that for a given metric space (X, d) of size n, with probability at least $\frac{2}{3}$ estimates $cost(G)$ to within a $1 + \epsilon$ factor requires $\Omega(n^2/k)$ time. Similarly, any algorithm that with probability at least $\frac{2}{3}$ estimates $cost(G)$ to within an additive error term $\epsilon(MST(X) + cost(X))$ requires $\Omega(\min\{nk^{3/2}, n^2/k\})$ time.

Artur Czumaj
Department of Computer Science and DIMAP
University of Warwick
A.Czumaj@warwick.ac.uk

Christian Sohler
Department of Mathematics and Computer Science
University of Cologne, Germany
csohler@uni-koeln.de

CP46

Improved Local Computation Algorithm for Set Cover Via Sparsification

We design a Local Computation Algorithm (LCA) for the set cover problem. Given a set system where each set has size at most s and each element is contained in at most t sets, the algorithm reports whether a given set is in some fixed set cover whose expected size is $O(\log s)$ times the minimum fractional set cover value. Our algorithm requires $s^{O(\log s)} t^{O(\log s \cdot (\log\log s + \log\log t))}$ queries. This result improves upon the application of the reduction of [Parnas and Ron, TCS'07] on the result of [Kuhn et al., SODA'06], which leads to a query complexity of $(st)^{O(\log s \cdot \log t)}$. To obtain this result, we design a parallel set cover algorithm that admits an efficient simulation in the LCA model by using a *sparsification* technique introduced in [Ghaffari and Uitto, SODA'19] for the maximal independent set problem. The parallel algorithm adds a random subset of the sets to the solution in a style similar to the PRAM algorithm of [Berger et al., FOCS'89]. However, our algorithm differs in the way that it never *revokes* its decisions, which results in a fewer number of adaptive rounds. This requires a novel approximation analysis which might be of independent interest.

Christoph Grunau
ETH Zurich
cgrunau@student.ethz.ch

Slobodan Mitrovic
MIT
slobo@mit.edu

Ronitt Rubinfeld
CSAIL, MIT and Tel Aviv University
ronitt@csail.mit.edu

Ali Vakilian
MIT
CSAIL
vakilian@mit.edu

CP46

Nearly Optimal Edge Estimation with Independent Set Queries

We study the problem of estimating the number of edges of an unknown, undirected graph $G = ([n], E)$ with access to an independent set oracle. When queried about a subset $S \subseteq [n]$ of vertices, the independent set oracle answers whether S is an independent set in G or not. Our first main result is an algorithm that computes a $(1 + \epsilon)$-approximation of the number of edges m of the graph using $\min(\sqrt{m}, n/\sqrt{m}) \cdot \text{poly}(\log n, 1/\epsilon)$ independent set queries. This improves the upper bound of $\min(\sqrt{m}, n^2/m) \cdot \text{poly}(\log n, 1/\epsilon)$ by Beame et al. [BHRRS18]. Our second main result shows that $\min(\sqrt{m}, n/\sqrt{m})/\text{polylog}(n)$ independent set queries are necessary, thus establishing that our algorithm is optimal up to a factor of $\text{poly}(\log n, 1/\epsilon)$.

Amit Levi
University of Waterloo
amit.levi@uwaterloo.ca

Xi Chen, Erik Waingarten
Columbia University
xichen@cs.columbia.edu., eaw@cs.columbia.edu.

CP46

Robust Clustering Oracle and Local Reconstructor of Cluster Structure of Graphs

We develop sublinear time algorithms for analyzing the cluster structure of graphs with noisy partial information. A graph G with maximum degree at most d is called to be an ϵ-perturbation of a (k, ϕ_i, ϕ_o)-clusterable graph, if there is a partition of G with at most k parts (called clusters), such that one can insert/delete at most ϵdn intra-cluster edges to make each part have inner conductance at least ϕ_i, and outer conductance at most ϕ_o. We assume d is constant and are given query access to the adjacency list of such a graph. We construct in $O(\sqrt{n}poly(\frac{k \log n}{\phi\epsilon}))$ time a robust clustering oracle for a bounded-degree graph G that is an ϵ-perturbation of a $(k, \phi, O(\frac{\epsilon\phi}{k^3 \log n}))$-clusterable graph. Using such an oracle, a typical clustering query (e.g., IsOutlier(s), SameCluster(s, t)) can be answered in $O(\sqrt{n}poly(\frac{k \log n}{\phi\epsilon}))$ time and the answers are consistent with a partition of G in which all but $O(k\sqrt{\epsilon/\phi}n)$ vertices belong to a good cluster, i.e., a set with inner conductance at least $\phi/2$, and outer conductance $O(\frac{\sqrt{\epsilon}\phi^{1.5}}{k^4 \log n}))$. We also give a local reconstructor with similar runtime guarantee that takes as input a graph as above, and on any query vertex v, outputs all its neighbors in the reconstructed graph G', which is clusterable and close to the original graph.

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

CP46

A Lower Bound on Cycle-Finding in Sparse Digraphs

We consider the problem of finding a cycle in a sparse directed graph G that is promised to be far from acyclic, meaning that the smallest feedback arc set in G is large. We prove an information-theoretic lower bound, showing that for N-vertex graphs with constant outdegree any algorithm for this problem must make $\tilde{\Omega}(N^{5/9})$ queries to an adjacency list representation of G. In the language of property testing, our result is an $\tilde{\Omega}(N^{5/9})$ lower bound on the query complexity of one-sided algorithms for testing whether sparse digraphs with constant outdegree are far from acyclic. This is the first improvement on the $\Omega(\sqrt{N})$ lower bound, implicit in Bender and Ron, which follows from a simple birthday paradox argument.

Timothy Sun, Xi Chen, Tim Randolph, Rocco A. Servedio
Columbia University
tim@cs.columbia.edu, xichen@cs.columbia.edu, t.randolph@columbia.edu, rocco@cs.columbia.edu

Speaker Index

siam 2020

ACM-SIAM Symposium on
Discrete Algorithms
January 5–8, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Algorithm Engineering and Experiments
January 5–6, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Simplicity in Algorithms
January 6–7, 2020 • Salt Lake City, Utah, U.S.

Symposium on
Algorithmic Principles of Computer Systems
January 8, 2020 • Salt Lake City, Utah, U.S.

Italicized names indicate session organizers

M

N

O

P

Q

R

S

T

U

V

W

X

Y

Z

Notes

Hilton Salt Lake City Center Floor Plans