siam 2023

Searchable Abstracts Document

Conference on Applied and Computational Discrete Algorithms

(ACDA23)

May 31–June 2, 2023

This document was current as of May 22, 2023. Abstracts appear as submitted.

siam

3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 U.S.
Telephone: 800-447-7426 (U.S. & Canada) +1-215-382-9800 (Worldwide)
meetings@siam.org

IP2

Bigger Is Better: Why Wafer-Size Chips Matter

To extend the growth of performance in HPC, architectural specialization has returned, providing a significant boost for computational science. At the hardware level, the development by Cerebras Systems of a successful wafer-scale compute platform opens new opportunities. Single-wafer integration provides remarkable memory and interconnect bandwidth and latency. Performance is attainable without relying on dense matrix multiplication; there is no cache hierarchy and fine-grained processor to processor communication is efficient. Therefore, wafer-scale processing has breakthrough implications across high-performance use cases. The elimination of most off-chip communication cuts the power per unit performance, a key performance determining parameter. For HPC use cases like finite difference methods, Fourier methods, dense matrices, and examples involving irregular and dynamic interprocessor communication, we have demonstrated supercomputer performance levels or better at kilowatt rather than megawatt power and single rack rather than datacenter cost. The fine-grained nature of the architecture allows strong scaling and reduces time to solution rather than boosting flops for enormous problems.

Rob Schreiber
Cerebras, U.S.
rob.schreiber@cerebras.net

IP3

Political Apportionment Through the Lens of Combinatorial Optimization

The question of how to allocate seats in a body of representatives is one of the most fundamental tasks in the political organization of modern societies. In the apportionment problem, we are given the population of each state and the house size, and the goal is to design a method to allocate a certain number of seats to each state. Two fundamental properties designed to avoid the famous Alabama paradox and population paradox are monotonicity and proportionality. A method is house-monotone if no state receives fewer seats when the house size increases, and is quota-compliant if each state receives the exact proportional value, rounded up or down. In this work, we provide a polyhedral characterization of all the possible apportionments that can be obtained by house-monotone and quota-compliant methods. Our result is based on modeling the seat allocation process as a network flow. Furthermore, our approach is flexible in the sense that we can incorporate any other apportionment requirements as long as the network structure is preserved. Since achieving exact proportionality is not possible in general, we use randomization in the design of our methods. A method is ex-ante proportional if, on expectation, every state receives the exact proportional amount. We design a method that is quota-compliant, ex-ante proportional, and satisfies the more general population monotonicity property at the cost of violating the house size by some amount. Furthermore, by combining our polyhedral characterization together with randomization, we are able to characterize the randomized apportionment solutions that are house-monotone, quota-compliant, and ex-ante proportional. This is joint work with Javier Cembrano, Alexandros Tsigonias-Dimitriadis, and Victor Verdugo.

Jose Correa
Universidad de Chile, Chile
correa@uchile.cl

IP4

New Local Differentially Private Protocols for Frequency and Mean Estimation

Consider the following examples of distributed applications: a texting app wants to train ML models for autocomplete based on text history residing on-device across millions of devices, or the developers of some other app want to understand common app settings by their users. In both cases, and many others, a third party wants to understand something in the aggregate about a large distributed database but under the constraint that each individual record requires some guarantee of privacy. Protocols satisfying so-called local differential privacy have become the gold standard for guaranteeing privacy in such situations, and in this talk I will discuss new such protocols for two of the most common problems that require solutions in this framework: frequency estimation, and mean estimation. Based on joint works with subsets of Hilal Asi, Vitaly Feldman, Huy Le Nguyen, and Kunal Talwar.

Jelani Nelson
University of California, Berkeley
minilek@berkeley.edu

IP5

Where the Wild Things Are

As computer scientists and applied mathematicians, we strive to develop novel algorithms, software, and systems that provide users with significantly expanded capabilities. In some cases, a single effort achieves both goals, but in many cases we must choose between innovation and utility. The academics among us are also responsible for training the next generation. The talk will describe a decade-long adventure that started as an effort to provide ecologists who track wild animals significantly expanded capabilities. The effort succeeded; the new tracking capabilities led to discoveries published in Science, PNAS, as well as more specialized journals. This application-oriented effort also resulted in opportunities to innovate in algorithms, systems, and mathematics, including a new mixed-integer least-squares algorithm for specialized GPS localization and a (correct) proof to a 60-year-old mathematical claim. The insights gained in the context of this research led to a new university course and to a new SIAM book, Location Estimation from the Ground Up. The talk will describe the most interesting aspects of this adventure, will describe remaining challenges and open problems, and will also reflect on different drivers for research and different ways of doing research.

Sivan A. Toledo
Tel Aviv University
stoledo@tau.ac.il

IP6

Not Too Sparse, Not Too Dense ...

Over the last ten years, there has been significant work on adapting tools and techniques from structural graph algorithms for practical network analysis. One promising direction is bounded expansion - a broad notion of sparsity introduced by Neetril and Ossona de Mendez in 2012 which seems to exhibit a Goldilocks-like property of both encompassing many classes of real-world networks while still ad-

mitting (theoretically) efficient parameterized algorithms for a rich set of graph optimization problems. Bounded expansion includes many well-studied classes (e.g. bounded degree, H-minor-free, bounded crossing number) and can be characterized by a plethora of independently interesting parameters (e.g. density of shallow minors, uniform quasi-wideness, and generalized coloring numbers). Many FPT algorithms take advantage of the inherent bounds on these parameters in bounded expansion classes in order to achieve polynomial runtimes. In this talk, we focus on the generalized coloring numbers, which were introduced by Kierstead and Yang in 2003 and have played a key role in work on practical implementations of algorithms in structurally sparse graphs. We survey the area, highlighting the use of generalized coloring numbers in engineering a new algorithm for counting subgraphs in large sparse networks (Algorithmica 2022, with Reidl) alongside recent work on their hardness and approximation (EuJC 2023, with Breen-McKay and Lavallee).

Blair D. Sullivan
University of Utah, U.S.
sullivan@cs.utah.edu

SP1

Session 1 - 2023 SIAG/ACDA Early Career Prize: Generalized Hypergraph Cut Algorithms and Their Applications

In graph clustering and hypergraph clustering, the goal is to partition a set of nodes into well-connected clusters in such a way that few edges (or hyperedges) are "cut", i.e., separated between clusters. Although graph models are more common, hypergraphs can directly encode multiway relationships like group social interactions, multiway biological interactions, or chemical reactions involving multiple substances. As a result, hypergraph clustering is more natural for many applications. However, this generalized problem is more nuanced since there are many different ways to partition a single hyperedge among clusters. Different ways of penalizing cut hyperedges may be more or less useful depending on the application, and this also leads to significant differences in computational complexity results when solving hypergraph clustering problems. This talk will introduce a generalized notion of the minimum hypergraph cut problem, along with efficient algorithms for certain variants and NP-hardness results for others. I'll then show how new algorithmic techniques for hypergraph cut problems can be incorporated into more sophisticated pipelines for hypergraph clustering. This leads to improved results in several different types of downstream data analysis tasks, such as finding related retail products in large e-commerce datasets, detecting groups of students from the same classroom based on group interaction data, and identifying related posts in large online QA forums.

Nate Veldt
Texas A & M University
nveldt@tamu.edu

JP1

Welcome Remarks and Joint Plenary with OP23 - On the Optimization of Nonsmooth Problem Without Generalized Derivatives

Numerous applications may require the solution of nonsmooth optimization problems, e.g., due to the action of values, a bilevel structure or complementarity conditions. For a large class of target functions, the concept of abs-linearization allows the generation of a piecewise linear local model that is provable of second order. Similar to the quadratic model generated by a truncated Taylor series in the smooth situation, this piecewise linear model can be used as a building block for optimization algorithms targeting nonsmooth problems of different kinds. In this talk, first we define abs-smooth functions covering a wide range of applications like clustering, image restoration, and robust gas network optimization. Also various mathematical models like complementarity problems or bilevel optimization tasks can be formulated as abs-smooth functions. Subsequently, the abs-linearization approach and the properties of the resulting local model will be illustrated. Then, we discuss the solution of piecewise linear optimization problems. Based on this capability we can derive optimization approaches that allow also the handling of constraints and/or nonlinearities. For each class of the nonsmooth optimization problems and the corresponding solution approach, convergence results and numerical examples will be given.

Andrea Walther
Humboldt-Universität zu Berlin
andrea.walther@math.hu-berlin.de

CP1

Exact Algorithms for Group Closeness Centrality

The Group Closeness Centrality problem asks, given a graph G and an integer k, for a vertex set S of size k such that the sum of distances from the vertices of $V \setminus S$ to S is minimal. Being a generalization of the NP-hard Dominating Set problem, Group Closeness Centrality is NP-hard as well and W[2]-hard with respect to k meaning that it presumably has no algorithm with running time $f(k) \cdot n^{\mathcal{O}(1)}$. We first show that, in contrast to Dominating Set, Group Closeness Centrality remains W[2]-hard when restricted to graphs with constant maximum degree. We then develop and evaluate two new exact algorithms for Group Closeness Centrality, one based on branch-and-bound and one based on a new ILP formulation. We further show how to embed both approaches in an iterative method that allows to solve the problem without computing all pairwise distances in G. Our experiments show that on small and medium-sized real-world networks, the new ILP formulation substantially outperforms a previous ILP formulation and that the branch-and-bound algorithm is competitive with the new ILP formulation for small values of k.

Luca P. Staus, Christian Komusiewicz, Nils Morawietz
Philipps-Universität Marburg, Germany
staus@students.uni-marburg.de,
komusiewicz@informatik.uni-marburg.de,
morawietz@informatik.uni-marburg.de

Frank Sommer
Philipps-Universität Marburg
fsommer@informatik.uni-marburg.de

CP1

On the Group Coverage Centrality Problem: Parameterized Complexity and Heuristics

We study the problem of computing a group of k vertices that covers a maximum number of shortest paths in a graph in the context of network centrality. In Group Coverage Centrality the input is an undirected graph $G = (V, E)$ and the aim is to find a set S of k vertices such that the number of vertex pairs $\{u, v\}$ covered by S is at least t,

if such a set exists. A set S covers a vertex pair $\{u, v\}$ if S contains at least one internal vertex of some shortest (u, v)-path. We also consider All Pairs Coverage, the special case of Group Coverage Centrality where we want to cover at least one shortest path for all non-adjacent vertex pairs. We provide a parameterized complexity analysis for both problems for the solution-size related parameters k and $|V| - k$, structural graph parameters, and the number of covered shortest paths t. On the negative side, we show that solution-size parameterizations do not lead to FPT-algorithms. On the positive side, we show that additionally considering structural parameters motivated from social network theory leads to FPT-algorithms for All Pairs Coverage and that Group Coverage Centrality admits FPT-algorithms for t and for $t - k$. On the practical side, we introduce several heuristics and compare their performance on standard benchmark graphs. We show that a greedy Set Cover-based heuristic gives almost-optimal results while a simple degree-based heuristic performs only slightly worse with a much better running time.

Christian Komusiewicz, Jannis Schramek, Frank Sommer
Philipps-Universität Marburg
komusiewicz@mathematik.uni-marburg.de,
schramek@students.uni-marburg.de,
fsommer@informatik.uni-marburg.de

CP2

Engineering Fully Dynamic Δ-Orientation Algorithms

In this work, we consider the fully dynamic edge orientation problem, also called fully dynamic Δ-orientation problem, which is to maintain an orientation of the edges of an undirected graph such that the out-degree is low. If edges are inserted or deleted, one may have to flip the orientation of some edges in order to avoid vertices having a large out-degree. While there has been theoretical work on dynamic versions of this problem, currently there is no experimental evaluation available. In this work, we close this gap and engineer a range of new dynamic edge orientation algorithms as well as algorithms from the current literature. The most successful newly proposed algorithms in this papers are based on repeatedly swapping edges with smallest degree neighbors and based on breadth-first search. Generally, these algorithms outperform the algorithms currently proposed in the literature. The best algorithm considered in this paper in terms of quality, based on a simple breadth-first search, computes the optimum result more than 90% of the instances and is on average only 2.4% worse than the optimum solution.

Jannick Borowitz, Ernestine Großmann
Heidelberg University
jannick.borowitz@uni-heidelberg.de,
e.grossmann@informatik.uni-heidelberg.de

Christian Schulz
Heidelberg University, Germany
christian.schulz@informatik.uni-heidelberg.de

CP2

On Improving Contraction Hierarchies Using Distance Group Partitions

Contraction Hierarchies (CH) are a popular speed-up technique for shortest path computations in transportation networks. One shortcoming of this approach is that it does not allow flexible trade-offs between query time and space consumption. We address this issue and show a way to spend additional space in order to improve the query times of CH. For instance, with only 5% additional space overhead we are able to improve the query times for the street network of South America by a factor of three. Our work is based on the notion of Distance Group Partitions, which we consider to be a natural extension of the CH framework.

Claudius Proissl
University of Stuttgart
claudius.proissl@fmi.uni-stuttgart.de

CP2

Faster Parallel Exact Density Peaks Clustering

Clustering multidimensional points is a fundamental data mining task, with applications in many fields. The goal of clustering algorithms is to group similar objects together. Density-based clustering is a clustering approach that defines clusters as dense regions of points. It has the advantage of being able to detect clusters of arbitrary shapes, rendering it useful in many applications. In this paper, we propose two fast parallel algorithms for Density Peaks Clustering (DPC), a popular variant of density-based clustering. Existing exact DPC algorithms suffer from low parallelism, which limits their application to large-scale data sets. Our most performant algorithm uses priority search kd-trees to compute DPC. It achieves O(log n log log n) span (parallel time complexity) while being work-efficient, which means it has the same work complexity as the best existing sequential DPC algorithm. We also present another DPC algorithm that uses a Fenwick tree and makes fewer assumptions for its average-case complexity to hold. We provide optimized implementations of our algorithms and evaluate their performance via extensive experiments. On a 30-core machine with two-way hyperthreading, our best algorithm achieves a 10.8–13169x speedup over the previous best parallel exact DPC algorithm. Compared to the state-of-the-art parallel approximate DPC algorithm, our best algorithm achieves a geometric mean speedup of 55.8x while being exact.

Yihao Huang
Phillips Academy
yhuang23@andover.edu

Shangdi Yu, Julian Shun
MIT
shangdiy@mit.edu, jshun@mit.edu

CP3

Integer Programming for the Maximum Cut Problem: A Refined Model and Implications for Branching

We suggest a refined integer programming model for the Maximum Cut problem, which is motivated by new and extended structural results. These include new insights on odd-cycle inequalities based on uncovering and making use of their relation to parity constraints, as well as implications of branching on partitioning decisions. The latter lead to a new improved branching rule. In our computational comparison of branching strategies, we demonstrate a significant impact of our model and techniques on state-of-the-art branch-and-cut algorithms and show that our branching rule improves over previous problem-specific

ones.

Jonas Charfreitag, Sven Mallach, Petra Mutzel
University of Bonn
j.charfreitag@gmail.com, sven.mallach@cs.uni-bonn.de, petra.mutzel@cs.uni-bonn.de

CP3

A Reduced Cost-Based Model Strengthening Method

Reduced cost strengthening is a common tool in Integer Linear Programming to speed up Branch-and-Bound algorithms. Given an optimal dual solution and a primal bound P, this technique allows to eliminate values from a variable's domain that cannot be attained in a solution with an objective value better than P. In this work, we present new theoretical results that allow for a different view of reduced costs. Based on these results, we designed a procedure that determines alternative dual solutions in order to strengthen the bound of any variable, i.e., it extends the classical reduced cost strengthening to variables that have only implicit bounds or are fractional. In the latter case, we even cut off the current fractional solution. This proves particularly useful in a Branch-and-Price context. We carry out computational experiments for a heterogeneous vehicle routing problem which show that our technique can have a large positive impact on the solution process.

Lukas Schürmann
University of Bonn
Institute for Computer Science
lschuermann@cs.uni-bonn.de

Petra Mutzel
University of Bonn
petra.mutzel@cs.uni-bonn.de

CP4

Invited Minitutorial 1 - Sergei Vassilvitskii; Beyond Worst Case Analysis

Worst-case analysis has proven invaluable for understanding aspects of both the complexity and practicality of algorithms. In some cases, however, we do not face worst-case scenarios, and the question arises of how we can tune our algorithms to work even better on the kinds of instances we are likely to see, while keeping a rigorous formal framework similar to what we have developed through worst-case analysis. We give an overview of a recent trend that develops algorithms parameterized by additional parameters which capture "the kinds of instances we are likely to see," and obtains a finer grained analysis of algorithms' performance. We will give examples of re-analyzing classical algorithms through this lens, as well as developing new algorithms that expose new structural insights about the problems.

Sergei Vassilvitskii
Google Research
sergei@cs.stanford.edu

CP5

Irreducibility of Recombination Markov Chains in the Triangular Lattice

In the United States, regions are frequently divided into districts for the purpose of electing representatives. How the districts are drawn can affect who's elected, and drawing districts to give an advantage to a certain group is known as gerrymandering. It can be surprisingly difficult to detect gerrymandering, but one algorithmic method is to compare a current districting plan to a large number of randomly sampled plans to see whether it is an outlier. Recombination Markov chains are often used for this random sampling: randomly choose two districts, consider their union, and split this union in a new way. This works well in practice, but the theory behind it remains underdeveloped. For example, it's not known if recombination Markov chains are irreducible, that is, if recombination moves suffice to move from any districting plan to any other. Irreducibility of recombination Markov chains can be formulated as a graph problem: for a graph G, is the space of all partitions of G into k connected subgraphs (k districts) connected by recombination moves? We consider three simply connected districts and district sizes $k_1 \pm 1$ vertices, $k_2 \pm 1$ vertices, and $k_3 \pm 1$ vertices. We prove for arbitrarily large triangular regions in the triangular lattice, recombination Markov chains are irreducible. This is the first proof of irreducibility under tight district size constraints for recombination Markov chains beyond small or trivial examples.

Sarah Cannon
Claremont McKenna College
scannon@cmc.edu

CP5

Generalized Scaling for the Constrained Maximum-Entropy Sampling Problem

The best practical techniques for exact solution of instances of the constrained maximum-entropy sampling problem, a discrete-optimization problem arising in the design of experiments, are via a branch-and-bound framework applied to a variety of concave continuous relaxations of the objective function. A standard and computationally-important bound-enhancement technique in this context is *scaling*, via a single positive parameter. We extend this technique to *generalized scaling*, employing a positive vector of parameters, we give mathematical results aimed at supporting algorithmic methods for computing optimal generalized scalings, and we give computational results demonstrating the usefulness of generalized scaling on benchmark problem instances.

Zhongzhu Chen, Jon Lee
University of Michigan
zhongzhc@umich.edu, jonxlee@umich.edu

Marcia Fampa
Federal University of Rio de Janeiro
fampa@cos.ufrj.br

CP5

Learned Interpolation for Better Streaming Quantile Approximation with Worst-Case Guarantees

An ε-approximate quantile sketch over a stream of n inputs approximates the rank of any query point $q$that is, the number of input points less than $q$up to an additive error of εn, generally with some probability of at least $1 - 1/\text{poly}(n)$, while consuming $o(n)$ space. While the celebrated KLL sketch of Karnin, Lang, and Liberty achieves a provably optimal quantile approximation algorithm over worst-case streams, the approximations it achieves in practice are often far from optimal. Indeed, the most commonly

used technique in practice is Dunning's t-digest, which often achieves much better approximations than KLL on real-world data but is known to have arbitrarily large errors in the worst case. We apply interpolation techniques to the streaming quantiles problem to achieve better approximations on real-world data sets than KLL while maintaining similar guarantees in the worst case.

Nicholas Schiefer
Anthropic
schiefer@mit.edu

Justin Y. Chen
MIT
justc@mit.edu

Piotr Indyk, Shyam Narayanan
Massachusetts Institute of Technology
indyk@mit.edu, shyamsn@mit.edu

Sandeep Silwal
Massachusetts Institute of Technology, U.S.
silwal@mit.edu

Tal Wagner
MIT
MIT
tal.wagner@gmail.com

CP6

Reducing the Feeder Effect in Public School Admissions: A Bias-Aware Analysis for Targeted Interventions

Traditionally, New York Citys top 8 public schools select candidates solely based on their scores in the Specialized High School Admissions Test (SHSAT). These scores are known to be impacted by socioeconomic status of students and test preparation received in feeder middle schools, leading to a massive filtering effect in the education pipeline. The classical mechanisms for assigning students to schools do not naturally address problems like school segregation and class diversity, which have worsened over the years. The scientific community, including policy makers, have reacted by incorporating group-specific quotas or summer school opportunities, but there is evidence that these can end up hurting minorities or even create legal challenges. We take a completely different approach in this work, to reduce this filtering effect of feeder middle schools. We model a two-sided market where a candidate may not be perceived at their true potential, and is therefore assigned to a lower level than the one (s)he deserves. We propose and study the effect of interventions such as additional training and scholarships towards disadvantaged students, and challenge existing mechanisms for scholarships. We validate these findings using SAT scores data from New York City high schools. We further show that our qualitative takeaways remain the same even when some of the modeling assumptions are relaxed.

Yuri Faenza
Columbia University
yf2414@columbia.edu

Swati Gupta
Georgia Institute of Technology, U.S.
swatig@gatech.edu

Xuan Zhang
Meta
xuanzh@meta.com

CP6

An Exponentially Smaller Kernel for Exact Weighted Clique Decomposition

Mining groups of genes that consistently co-express is an important problem in biomedical research. Recently, Cooley et al. modeled this problem as Exact Weighted Clique Decomposition (EWCD) in which, given an edge-weighted graph G and a positive integer k, the goal is to decompose G into at most k (overlapping) weighted cliques so that an edge's weight is exactly equal to the sum of weights for cliques it participates in. They show that EWCD is fixed-parameter-tractable, giving a 4^k-kernel alongside a backtracking algorithm to iteratively build a decomposition. Unfortunately, because of inherent exponential growth in the space of potential solutions, their method is typically able to decompose graphs only when $k \leq 11$. In this work, we establish reduction rules that exponentially decrease the size of the kernel (from 4^k to $k2^k$) and use insights about the structure of potential solutions to give new search rules that speed up the decomposition algorithm. At the core of our techniques is a result from combinatorial design theory called Fisher's inequality characterizing set systems with restricted intersections. Experimental evaluation of our kernelization and decomposition algorithms (together called DeCAF) on a corpus of biologically-inspired data showed that in most cases DeCAF leads to over 80% reduction in the size of the kernel, orders of magnitude improvement in the time required to obtain a decomposition, and scales to instances with $k \geq 17$.

Shweta Jain, Yosuke Mizutani
University of Utah
shwetaj@illinois.edu, yos@cs.utah.edu

Blair D. Sullivan
University of Utah, U.S.
sullivan@cs.utah.edu

CP6

Order-Optimal Correlated Rounding for Fulfilling Multi-Item E-Commerce Orders

We study the dynamic fulfillment problem in e-commerce, in which incoming (multi-item) customer orders must be immediately dispatched to (a combination of) fulfillment centers that have the required inventory. A prevailing approach to this problem, pioneered by Jasin and Sinha (2015), is to write a "deterministic' linear program that dictates, for each item in an incoming multi-item order from a particular region, how frequently it should be dispatched to each fulfillment center (FC). However, dispatching items in a way that satisfies these frequency constraints, without splitting the order across too many FC's, is challenging. Jasin and Sinha identify this as a *correlated rounding* problem, and propose an intricate rounding scheme that they prove is suboptimal by a factor of at most $\approx q/4$ on a q-item order. This paper provides to our knowledge the first substantially improved scheme for this correlated rounding problem, which is suboptimal by a factor of at most $1 + \ln(q)$. We provide another scheme for sparse networks, which is suboptimal by a factor of at most d if each item is stored in at most d FC's. We show both of these guarantees to be tight in terms of the dependence on q or d. Our schemes are simple and fast, based on an intuitive idea—items wait for FC's to "open' at random times, but

observe them on "dilated' time scales. This also implies a new randomized rounding method for the classical Set Cover problem.

Will Ma
Columbia University
wm2428@gsb.columbia.edu

CP7

Provably Fast and Space-Efficient Parallel Biconnectivity

Computing biconnected components (BCC) of a graph is a fundamental graph problem. The canonical parallel BCC algorithm is the Tarjan-Vishkin algorithm, which has $O(n+m)$ optimal work and polylogarithmic span on a graph with n vertices and m edges. However, Tarjan-Vishkin is not widely used in practice. We believe the reason is the space-inefficiency (it uses $O(m)$ extra space). In practice, existing parallel implementations are based on breath-first search (BFS). Since BFS has span proportional to the diameter of the graph, existing parallel BCC implementations suffer from poor performance on large-diameter graphs and can be slower than the sequential algorithm on many real-world graphs. We propose the first parallel biconnectivity algorithm (FAST-BCC) that has optimal work, polylogarithmic span, and is space-efficient. Our algorithm creates a skeleton graph based on any spanning tree of the input graph. Then we use the connectivity information of the skeleton to compute the biconnectivity of the original input. We carefully analyze the correctness of our algorithm, which is highly non-trivial. We implemented FAST-BCC and compared it with existing implementations. We tested them on a 96-core machine on 27 graphs with varying edge distributions. FAST-BCC is the fastest on *all* graphs. On average (geometric means), FAST-BCC is 3.1× faster than the *best existing baseline* on each graph.

Xiaojun Dong, Letong Wang, Yan Gu
University of California, Riverside
xdong038@ucr.edu, lwang323@ucr.edu, ygu@cs.ucr.edu

Yihan Sun
UC Riverside
yihans@cs.ucr.edu

CP7

Jet: Multilevel Graph Partitioning on Graphics Processing Units

The multilevel heuristic is the dominant strategy for high-quality sequential and parallel graph partitioning. Partition refinement is a key step of multilevel graph partitioning. In this work, we present Jet, a new parallel algorithm for partition refinement specifically designed for Graphics Processing Units (GPUs). We combine Jet with GPU-aware coarsening to develop a k-way graph partitioner, the Jet partitioner. The new partitioner achieves superior quality when compared to state-of-the-art shared memory partitioners on a large collection of test graphs.

Michael Gilbert
Pennsylvania State University, U.S.
msg5334@psu.edu

Kamesh Madduri
Pennsylvania State University
madduri@psu.edu

Erik G. Boman
Center for Computing Research
Sandia National Labs
egboman@sandia.gov

Sivasankaran Rajamanickam
Sandia National Laboratories
srajama@sandia.gov

CP8

Computing Boundary Crossing Probabilities of General Empirical Processes

Order statistics play a fundamental role in statistical procedures such as risk estimation, outlier detection, and multiple hypothesis testing, as well as in the analyses of mechanism design, queues, load balancing, and rank-based processes. They are particularly useful in the construction of empirical distribution functions, which are fundamental to empirical process theory and non-parametric statistics. In some applications, it may be desirable to compute the joint cumulative distribution function (cdf) of d order statistics *exactly*. Efficient algorithms to compute this quantity are known when the data are i.i.d.; however, the task becomes significantly more challenging when relaxing either the identically distributed or the independence assumption. Existing methods for the non-i.i.d. setting obtain the joint cdf indirectly, by first computing and then aggregating over the marginal distributions. In this paper, we simplify an existing dynamic programming solution to achieve an exponential-in-d factor improvement in both time and space complexity over known methods. We detail the independent, non-identical setting, and then outline how our method extends to more general settings: e.g., dependent random variables. (Details are deferred to an appendix.)

Rigel C. Galgana, Amy Greenwald, Takehiro Oyakawa
Brown University
galganarigel@gmail.com, amy@cs.brown.edu, takehiro_oyakawa@alumni.brown.edu

CP8

Greed Is Good for 2-Approximating the Jump Number for Interval Posets

The jump number problem for interval posets is NP-hard [Pulleyblank, On minimizing setups in precedence constrained scheduling, 1986]. Interval posets are a well-studied class of posets with applications in processor scheduling, where the jump number problem is one of the most important optimization problems. The *plain greed* algorithm refers to scheduling from start to finish, simply avoiding a jump unless there is no other option. This strategy when applied to general posets can be far from optimal ($n/2$ approximation ratio), but its effectiveness when applied to interval posets has been a matter of interest for decades [Felsner, 1990; Faigle and Schrader, 1985]. We show that plain greed achieves an approximation ratio of 2 for the jump number of interval posets; this is the first time an approximation ratio has been proved for plain greed on this class. Furthermore, the bound is tight, and we give an infinite class of interval posets for which greed yields exactly twice the optimal jump number. This resolves a decades old question about the effectiveness of greed for the jump number of interval posets. We introduce new techniques for establishing exact and approximate jump number strategies which may be applicable to other poset

problems.

Gara Pruesse
Vancouver Island University
gara.pruesse@viu.ca

CP9

LAHypergraph: Parallel Hypergraph Analytics in the Language of Linear Algebra

To design efficient and portable algorithms for hypergraphs, we consider the adjacency matrix and the incidence matrix representations of a hypergraph. An adjacency matrix (clique-expansion)-based hypergraph algorithm formulation has a one-to-one correspondence with processing a unipartite graph. However, in addition to large memory footprint, the adjacency matrix representation loses structural information about the original hypergraph. In contrast, an incidence matrix-based hypergraph algorithm formulation operates on a bipartite graph view of the modeled dataset. The incidence matrix representation retains both the hyperedge set and the vertex set information of a hypergraph and has lower memory footprint. Considering these facts, in this paper, we propose a suite of parallel, portable hypergraph algorithms, composed with a set of sparse linear algebra-based operations, especially for incidence matrix-based hypergraph processing. We identify the semiring operations for generalized sparse matrix-vector multiplication (SpMV) and the pattern of their applications, which are the main foundations of these algorithms. Implementations of our linear algebra based hypergraph algorithms for both the CPUs and the GPUs are included in a library, namely LAHypergraph. From a performance viewpoint, we demonstrate that our incidence matrix based algorithms in LAHypergraph outperform the state-of-the-art (SOTA) Hygra framework.

Luanzheng Guo
Pacific Northwest National Laboratory
lenny.guo@pnnl.gov

Jesun S. Firoz
Pacific Northwest National Lab
jesun.firoz@pnnl.gov

Gokcen Kestor
Pacific Northwest National Laboratory
gokcen.kestor@pnnl.gov

CP9

Solving Max-plus Linear Systems by Level Sparsification

In this study, an algorithm for solving max-plus linear systems is discussed. Such a system is applied to many problems, such as control and scheduling, and is equivalent to the mean-payoff game. To reduce the computational time required to solve max-plus linear systems, we improve the alternating method, which is a well-known algorithm for the problem. In max-plus arithmetics, small numbers rarely contribute to the results. Based on this observation, our method, called level sparsification, selects some large entries of matrices as active ones and ignore other entries by replacing with a fixed value. Depending on the number of selected entries, we propose two types of sparsification and experiment with the performance compared to the original alternating method.

Yuki Nishida
Doshisha University
ynishida@rs.tus.ac.jp

CP9

Accelerating Distributed Matrix Multiplication with 4-Dimensional Polynomial Codes

A single straggler worker may delay an entire distributed system. The state-of-the-art strategies for mitigating delays in large-scale distributed matrix multiplication are polynomial-based coded computations such as the Polynomial Codes and Entangled Polynomial Codes. While such strategies deal with stragglers efficiently, they discard partial computations performed by stragglers. Hence, they are sub-optimal. Here, we present the Multi Entangled Polynomial Codes, a straggler mitigation strategy that utilizes the computations performed by all workers and significantly reduces the running time. Furthermore, it allows the final output to be decoded before any worker completes its tasks, thereby breaking the lower bound of Yu, Maddah-Ali, and Avestimehr (2020). Previous studies that utilize partial computations performed by stragglers require large Maximal Distance Separable codes, resulting in high overhead costs. In contrast, our strategy requires short codes comparable to Entangled Polynomial Codes. Thus, we preserve efficient encoding and decoding complexity and reduce the arithmetic overhead of previous solutions by a factor of $O\left(N/W^{2/3}\right)$, where N and W are the matrices dimension and the number of workers, respectively. We provide experimental results on an Amazon EC2 cluster that demonstrate up to 15% speedup over previous strategies. Moreover, we show that our strategy is optimal up to a factor of $(1 + o(1))$.

Roy Nissim
The Hebrew University of Jerusalem
roynis1@gmail.com

Oded Schwartz
Hebrew University of Jerusalem
odedsc@cs.huji.ac.il

CP10

A Breakpoints Based Method for the Maximum Diversity and Dispersion Problems

The maximum diversity, or dispersion, problem (MDP), is to select from a given set a subset of elements of given cardinality (budget), so that the sum of pairwise distances, or utilities, between the selected elements is maximized. We introduce here a method, called the **Breakpoints** (BP) algorithm, based on a technique proposed in Hochbaum (2009), to generate the concave piecewise linear envelope of the solutions to the relaxation of the problem for all values of the budget. The breakpoints in this envelope are provably optimal for the respective budgets and are attained using a parametric cut procedure that is very efficient. The performance of the parametric cut is further enhanced by a newly introduced compact formulation of the problem. The problem is then solved, for any given value of the budget, by applying a greedy-like method to add or subtract nodes from adjacent breakpoints. This greedy solution may be improved using Tabu Search. the BPTS algorithm. This method works well if for the given budget there are breakpoints that are "close". If that is not the case, we introduce a perturbation technique that increases the number of breakpoints. We compare the performance of our breakpoints algorithm to leading methods for MDP

and demonstrate that our method dominates, dramatically, the performance of these methods. This is the first time that the properties of the concave envelope and the breakpoints are used in a practically tested algorithm.

Dorit Hochbaum
Industrial Engineering and Operations Research
University of California, Berkeley
hochbaum@ieor.berkeley.edu

Zhihao Liu
UC Berkeley
zhihao_liu@berkeley.edu

Olivier Goldschmidt
Riverside county office of education
goldoliv@gmail.com

CP10

Exponential Convergence of Sinkhorn Under Regularization Scheduling

In 2013, Cuturi [Cut13] introduced the Sinkhorn algorithm for matrix scaling as a method to compute solutions to regularized optimal transport problems. In this paper, aiming at a better convergence rate for a high accuracy solution, we work on understanding the Sinkhorn algorithm under regularization scheduling, and thus modify it with a mechanism that adaptively doubles the regularization parameter η periodically. We prove that such modified version of Sinkhorn has an exponential convergence rate as iteration complexity depending on $\log(1/\varepsilon)$ instead of $\varepsilon^{-O(1)}$ from previous analyses [Cut13] [altschuler2017near] in the optimal transport problems with integral supply and demand. Furthermore, with cost and capacity scaling procedures, the general *optimal transport* problem can be solved with a logarithmic dependence on $1/\varepsilon$ as well.

Jingbang Chen
University of Waterloo
chenjb1997@gmail.com

Li Chen
Georgia Institute of Technology
lichen@gatech.edu

Yang P. Liu
Stanford University
yangpliu@stanford.edu

Richard Peng
Georgia Institute of Technology, U.S.
rpeng@cc.gatech.edu

Arvind Ramaswami
Georgia Institute of Technology
aramaswami32@gatech.edu

CP10

Fast First-Order Methods for Monotone Strongly DR-Submodular Maximization

In this paper, we first introduce and characterize the class of strongly DR-submodular functions and show how such a property implies strong concavity along non-negative directions. Then, we study L-smooth monotone strongly DR-submodular functions that have bounded curvature, and we show how to exploit such additional structure to obtain algorithms with improved approximation guarantees and faster convergence rates for the maximization problem. In particular, we propose the SDRFW algorithm that matches the provably optimal $1 - \frac{c}{e}$ approximation ratio after only $\lceil \frac{L}{\mu} \rceil$ iterations, where $c \in [0, 1]$ and $\mu \geq 0$ are the curvature and the strong DR-submodularity parameter. Furthermore, we study the Projected Gradient Ascent (PGA) method for this problem and provide a refined analysis of the algorithm with an improved $\frac{1}{1+c}$ approximation ratio (compared to $\frac{1}{2}$ in prior works) and a linear convergence rate. Given that both algorithms require knowledge of the smoothness parameter L, we provide a novel characterization of L for DR-submodular functions showing that in many cases, computing L could be formulated as a convex optimization problem, i.e., a geometric program, that could be solved efficiently. Experimental results illustrate and validate the efficiency and effectiveness of our algorithms.

Omid Sadeghi, Maryam Fazel
University of Washington
omids@uw.edu, mfazel@uw.edu

CP11

Small Searchable K-Spectra Via Subset Rank Queries on the Spectral Burrows-Wheeler Transform

The k-spectrum of a string is the set of all distinct substrings of length k occurring in the string. This is a lossy but computationally convenient representation of the information in the string, with many applications in high-throughput bioinformatics. In this work, we define the notion of the Spectral Burrows-Wheeler Transform (SBWT), which is a particular sequence of subsets of the alphabet of the string that encodes k-spectrum of the string. The SBWT is a distillation of the ideas found in the BOSS and Wheeler graph data structures. We explore multiple different approaches to index the SBWT for membership queries on the underlying k-spectrum. We identify *subset rank queries* as the essential subproblem, and propose three space-efficient index structures to solve it. One of the approaches essentially leads to the known BOSS data structure, while the others offer attractive time-space trade-offs and support simpler query algorithms that rely only on fast rank queries. All of the approaches are also amendable to entropy compression, which leads to good space bounds on data structure size. We show, via experiments on a range of genomic data sets, that the simplicity of our new indexes translates into large performance gains in practice over prior art.

Jarno Alanko
University of Helsinki
jarno.alanko@helsinki.fi

Simon J. Puglisi
Department of Computer Science
University of Helsinki
simon.j.puglisi@gmail.com

Jaakko Vuohtoniemi
University of Helsinki
jaako.vuohtoniemi@helsinki.fi

CP11

Singleton Sieving: Overcoming the Memory/Speed

Trade-Off in Exascale K-Mer Analysis

Traditional filter data structures, such as Bloom filters, do not offer necessary features that modern high-performance data analytics applications need in order to efficiently perform complex data analysis tasks. For example, MetaHipMer, a de novo metagenome assembler, can use filters to weed out singleton k-mers and reduce memory usage by 30%-70%. However, the filter needs the ability to associate values with k-mers in order to perform the analysis in a single communication pass. Bloom filters do not support value associations and cause the application to perform an extra communication pass, thereby increasing the run time. Therefore, MetaHipMer faces a trade off between memory and speed due to the limited capabilities of traditional filters. In this paper, we overcome the memory and speed trade off in MetaHipMer by integrating a GPU-based feature-rich filter, the Two-Choice filter (TCF), in the MetaHipMer pipeline. The TCF uses key-value association to approximately store k-mers with extensions. This allows MetaHipMer to perform k-mer analysis on the GPUs in a single communication pass. Our empirical analysis shows a 50% reduction in memory usage in k-mer analysis on each node in MetaHipMer without any effect on the overall run time or assembly quality. The memory reduction in turn results in a 43% reduction in the number of nodes required to assemble datasets and enables MetaHipMer to scale to much larger datasets.

Hunter Mccoy
University of California, Berkeley
hunterm@berkeley.edu

Steven Hofmeyr
Lawrence Berkeley National Laboratory
shofmeyr@lbl.gov

Katherine Yelick
Lawrence Berkeley National Laboratory
University of California at Berkeley
yelick@eecs.berkeley.edu

Prashant Pandey
University of Utah
prashant.prashn@gmail.com

CP11

Elimination Techniques for Algorithmic Differentiation Revisited

All known elimination techniques for (first-order) algorithmic differentiation (AD) rely on Jacobians to be given for a set of relevant elemental functions. Realistically, elemental tangents and adjoints are given instead. They can be obtained by applying software tools for AD to the parts of a given modular numerical simulation. The novel generalized face elimination rule proposed in this article facilitates the rigorous exploitation of associativity of the chain rule of differentiation at arbitrary levels of granularity ranging from elemental scalar (state of the art) to multivariate vector functions with given elemental tangents and adjoints. The implied combinatorial Generalized Face Elimination problem asks for a face elimination sequence of minimal computational cost. Simple branch and bound and greedy heuristic methods are employed as a baseline for further research into more powerful algorithms motivated by promising first test results. The latter can be reproduced with the help of an open-source reference implementation.

Uwe Naumann, Erik Schneidereit, Simon Märtens,
Markus Towara
RWTH Aachen University, Germany
naumann@stce.rwth-aachen.de, schneidereit@stce.rwth-aachen.de, maertens@stce.rwth-aachen.de, towara@stce.rwth-aachen.de

PP1

Extending the Primal-Dual Method for Edge Connectivity

Input your abstract, including TeX commands, here. Williamson, Goemans, Vazirani and Mihail, in 1995, proposed the primal-dual method for edge augmentation problems where the connectivity requirements are given by the so called uncrossable functions. We address the long-standing open question of extending primal-dual methods to more general classes of functions. We see applications of our results in the recently studied flexible graph connectivity problem where we provide O(1) approximation algorithms. Applications are also seen in the capacitated k-edge connected problem.

Ishan Bansal, Logan Grout
Cornell University
ib332@cornell.edu, lcg58@cornell.edu

Sharat Ibrahimpur
London School of Economics
s.ibrahimpur@lse.ac.uk

Joseph Cheriyan
Department of Combinatorics and Optimization
University of Waterloo
jcheriyan@uwaterloo.ca

PP1

Vertex Elimination for Algorithmic Piecewise Differentiation

A piecewise affine approximation of a mathematical function can be represented in abs-normal form. Given a computer program that implements a mathematical function, this representation can be efficiently evaluated by algorithmic piecewise differentiation. By successively eliminating switching variables from such a representation, this work introduces the concept of reduced models. The novel contribution of this work is to describe the underlying phenomena from scientific computing by introducing a directed acyclic graph called condensed computational graph with corresponding vertex elimination sequences. This graph-theoretical approach allows the formulation of novel combinatorial optimization problems. For these problems defined on the condensed computational graph, new combinations of three quantitative measures are proposed: minimizing the number of floating-point operations, minimizing the length of the longest path, and minimizing the fill-in of a vertex elimination sequence. This goal-oriented approach paves the way not only for the design of certain properties of reduced models but also for the development of efficient graph-theoretical strategies for their construction.

Torsten F. Bosse
Friedrich-Schiller-Universität Jena
torsten.bosse@uni-jena.de

Martin Buecker
Mathematics and Computer Science Department
Friedrich Schiller University Jena

martin.buecker@uni-jena.de

PP1

Overlapping Edge-Colored Clustering

A recent trend in data mining has explored clustering algorithms for datasets with categorical relationship types. Such algorithms have applications in the analysis of social, co-authorship, and protein interaction networks, to name a few. Many such applications, however, naturally have some overlap between clusters, a nuance which is missing from current combinatorial models. We generalize a recent categorical clustering framework, Edge-Colored Clustering, to allow for a budgeted number of overlapping cluster assignments. We present algorithms which provide approximately minimal edge deletion sets, as well as bi-criteria approximations, where the second approximation factor is on the budget for cluster overlap.

Alex Crane, Brian Lavallee, Blair D. Sullivan
University of Utah
alex.crane@utah.edu, brian.lavallee@utah.edu, sullivan@cs.utah.edu

Nate Veldt
Texas A & M University
nveldt@tamu.edu

PP1

Practical Weighted Clique Decomposition for Identifying Gene Co-Expression Modules

Mining the membership of genes in modules and the relative strength of effect a module has on co-expression of its constituents is an important problem in biomedical research. Cooley et al. modeled this problem as Exact Weighted Clique Decomposition (EWCD) in which, given an edge-weighted graph G and an integer k, the goal is to decompose G into at most k (potentially overlapping) weighted cliques so that an edge's weight is exactly equal to the sum of weights for cliques it participates in. We propose an optimization variant of the problem which will produce a clique decomposition that minimizes its *discrepancy* with the original edge weights. In this work, we view the problem as a variant of Non-negative Matrix Factorization (NMF). Given a symmetric co-expression matrix, we aim to find a set of k cliques and their weights that minimize a specific loss function. We test our approaches with a comprehensive evaluation and comparison of our EWCD-based approach with outputs from Clust and WGCNA on several real-world gene networks, focusing on valudation using known modules.

Alex Crane, David Dursteler, Yosuke Mizutani, Blair D. Sullivan
University of Utah
alex.crane@utah.edu, u1161522@utah.edu, yos@cs.utah.edu, sullivan@cs.utah.edu

PP1

Counting Induced 6-Cycles in Bipartite Graphs

Various complex networks in real-world applications are best represented as a bipartite graph, such as user-product, paper-author, and actor-movie relations. Motif-based analysis has substantial benefits for networks and bipartite graphs are no exception. The smallest non-trivial subgraph in a bipartite graph is a (2,2)-biclique, also known as a butterfly. Although butterflies are succinct, they are limited in capturing the higher-order relations between more than two nodes from the same node set. One promising structure in this context is the induced 6-cycle which consists of three nodes on each node set forming a cycle where each node has exactly two edges. In this paper, we study the problem of counting induced 6-cycles through parallel algorithms. To the best of our knowledge, this is the first study on induced 6-cycle counting. We first consider two adaptations based on previous works for cycle counting in bipartite networks. Then, we introduce a new approach based on the node triplets and offer a systematic way to count the induced 6-cycles. Our final algorithm, BatchTripletJoin, is parallelizable across root nodes and uses minimal global storage to save memory. Our experimental evaluation on a 52 core machine shows that BatchTripletJoin is significantly faster than the other algorithms while being scalable to large graph sizes and number of cores. On a network with 112M edges, BatchTripletJoin is able to finish the computation in 78 mins by using 52 threads.

Jason Niu
SUNY University at Buffalo
jasonniu@buffalo.edu

Jaroslaw Zola, Ahmet Erdem Sariyüce
University at Buffalo
jzola@buffalo.edu, erdem@buffalo.edu

PP1

An Approximation of the Modulus of the Family of Edge Covers

The discrete p-modulus (modulus for short) is a very flexible and general tool for measuring the richness of families of objects defined on a graph. Modulus has been studied for specific families of graph objects and it has been shown to generalize well-known graph theoretic quantities such as shortest path, max flow/min cut, and effective resistance. Our focus is on the p-modulus of the family of edge covers and on the more general family of fractional edge covers on an unweighted, undirected graph. Through the theory of Fulkerson blocking duality, every family of objects has a corresponding dual family whose modulus is closely related to the modulus of the original family. Our results show that the dual family of fractional edge covers is the family of stars, which greatly reduces the number of constraints for the p-modulus problem. With this, we give an approximation for the modulus of edge covers using the modulus of fractional edge covers.

Adriana M. Ortiz Aquino
Kansas State University
Department of Mathematics
adrianaortiz@ksu.edu

Nathan Albin
Kansas State University
albin@ksu.edu

PP1

IcebergHT: High Performance Pmem Hash Tables Through Stability and Low Associativity

TBD

Prashant Pandey
University of Utah
prashant.prashn@gmail.com

Michael A. Bender
Stony Brook University
bender@cs.stonybrook.edu

Alex Conway
VMware Research Group
aconway@vmware.com

Martin Farach-Colton
Rutgers University
farach@cs.rutgers.edu

William Kuszmaul
Massachusetts Institute of Technology
kuszmaul@mit.edu

Guido Tagliavini
Rutgers University
guido.tag@rutgers.edu

Rob Johnson
VMWare Research
robj@vmware.com

PP1

Improved Space Bounds for Learning with Experts

We give improved tradeoffs between space and regret for the online learning with expert advice problem over T days with n experts. Given a space budget of n^δ for $\delta \in (0,1)$, we provide an algorithm achieving regret $\tilde{O}(n^2 T^{1/(1+\delta)})$, improving upon the regret bound $\tilde{O}(n^2 T^{2/(2+\delta)})$ in the recent work of Peng and Zhang [SODA23]. The improvement is particularly salient in the regime $\delta \to 1$ where the regret of our algorithm approaches $\tilde{O}_n(\sqrt{T})$, matching the T dependence in the standard online setting without space restrictions.

Anders Aamand
Massachusetts Institute of Technology, U.S.
aamand@mit.edu

Justin Y. Chen
MIT
justc@mit.edu

Huy Q. Nguyen
Princeton University
qn@princeton.edu

Sandeep Silwal
Massachusetts Institute of Technology, U.S.
silwal@mit.edu

PP1

Parallel Strong Connectivity Based on Faster Reachability

Computing strongly connected components (SCC) is among the most fundamental problems in graph processing. It is challenging in the parallel setting and is particularly hard on large-diameter graphs. Many existing parallel SCC implementations can be even slower than Tarjans sequential algorithm on large-diameter graphs. To tackle this challenge, we propose an efficient parallel SCC implementation using a new parallel reachability algorithm. Our solution is based on a novel idea referred to as vertical granularity control (VGC). It breaks the synchronization barriers to increase parallelism and hide scheduling overhead. To use VGC in our SCC algorithm, we also design an efficient data structure called the parallel hash bag. It uses parallel dynamic resizing to avoid redundant work in maintaining frontiers (vertices processed in a round). We implement the parallel BGSS SCC algorithm using our new parallel reachability algorithm. We compare our implementation to the state-of-the-art systems, including GBBS, iSpan, Multi-step, and our highly optimized Tarjans (sequential) algorithm, on 18 graphs, including social, web, k-NN, and lattice graphs. On a machine with 96 cores, our implementation is the fastest on 16 out of 18 graphs. On average over all graphs, our SCC is 6.0× faster than the best previous parallel code (GBBS), 12.8× faster than Tarjans sequential algorithms, and 2.7× faster than the best existing implementation on each graph.

Letong Wang, Xiaojun Dong, Yan Gu
University of California, Riverside
lwang323@ucr.edu, xdong038@ucr.edu, ygu@cs.ucr.edu

Yihan Sun
UC Riverside
yihans@cs.ucr.edu