

Breakthroughs in Theoretical Computer Science 2013, IIT Guwahati

Tuesday, December 10, 2013

08:30 - 09:25 **Registration**

09:30 - 10:30 **Costis Daskalakis (MIT)** – *Reductions from Mechanism to Algorithm Design*

10:30 - 11:00 **Tea/coffee break**

11:00 - 12:00 **Jelani Nelson (Harvard)** – *Dimensionality reduction via sparse matrices*

12:00 - 13:30 **Lunch break**

13:30 - 14:30 **Lap Chi Lau (CUHK)** - *Improved Cheeger's Inequality: Analysis of Spectral Partitioning Algorithms through Higher Order Spectral Gap*

14:30 - 15:30 **Lorenzo Orecchia (MIT)** – *A Simple, Combinatorial Algorithm for Solving SDD Systems in Nearly-Linear Time*

15:30 - 16:00 **Tea/coffee break**

16:00 - 17:00 **Thomas Rothvoss (MIT)** - *Better Bin Packing Approximations using Discrepancy Theory*

Wednesday, December 11, 2013

09:30 – 10:30 **Samuel Fiorini (ULB)** – *Exponential lower bounds for polytopes in combinatorial optimization*

10:30 - 11:00 **Tea/coffee break**

11:00 - 12:00 **Thomas Rothvoss (MIT)** - *The matching polytope has exponential extension complexity*

12:00 - 13:30 **Lunch break**

13:30 - 14:30 **Lorenzo Orecchia (MIT)** - *An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and its Multicommodity Generalizations*

14:30 - 15:30 **Nisheeth Vishnoi (MSR)** – *An Interior Point Method for Directed Max-Flow*

15:30 - 16:00 **Tea/coffee break**

16:00 - 17:00 **David Steurer (Cornell)** – *Sum-of-Squares and the Unique Games Conjecture*

17:00 - 18:00 **David Steurer (Cornell)** – *Analytic approach to Parallel Repetition*

Related Invited Talks at FSTTCS 2013

Friday, December 13, 2013. 09:00-10:00 **Venkat Guruswami (CMU)** - *Polar codes: Reliable communication with complexity polynomial in the gap to Shannon capacity*

Saturday, December 14, 2013. 09:00-10:00 **Subhash Khot (NYU)** - *On Approximation Resistance of Predicates*

Abstracts

Speaker: Costis Daskalakis

Title: Reductions from Mechanism to Algorithm Design

Abstract: Algorithmic mechanism design centers around the following question: How much harder is optimizing an objective over inputs that are furnished by rational agents compared to when the inputs are known? We present computationally efficient reductions from mechanism design (i.e. optimizing over rational inputs) to algorithm design (i.e. optimizing over known inputs) in general Bayesian settings. We also explore whether structural properties about optimal mechanisms can be inferred from these reductions. As an application, we present extensions of Myerson's celebrated single-item auction to multi-item settings.

Speaker: Samuel Fiorini

Title: Exponential lower bounds for polytopes in combinatorial optimization

Abstract: Linear programming is a powerful and widely used tool for tackling problems throughout Computer science and Mathematics. Despite the importance of linear programming, both in practice and theory, we do not have a good understanding of what problems can be efficiently expressed as linear programs (LPs). I will lift part of the veil and explain how to prove that there exist problems, such as the traveling salesperson problem (TSP), that need exponential size LPs. More precisely, our result states that every polytope that projects to the TSP polytope is defined by an exponential number of linear constraints. This result solves a 20-year old problem posed by Yannakakis. It was discovered through a new connection that we made between one-way quantum communication protocols and semidefinite programming reformulations of LPs. This is joint work with S. Massar (Brussels), S. Pokutta (Georgia Tech), H.R. Tiwary (Brussels), R. de Wolf (CWI).

Speaker: Lap Chi Lau

Title: Improved Cheeger's Inequality: Analysis of Spectral Partitioning Algorithms through Higher Order Spectral Gap

Abstract: Let $\phi(G)$ be the minimum conductance of an undirected graph G , and let $0 = \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n \leq 2$ be the eigenvalues of the normalized Laplacian matrix of G . We prove that for any graph G and any $k \geq 2$, $\phi(G) = O(k \lambda_2 / \sqrt{\lambda_k})$, and this performance guarantee is achieved by the spectral partitioning algorithm. This improves Cheeger's inequality, and the bound is optimal up to a constant factor for any k . The result shows that the spectral partitioning algorithm is a constant factor approximation algorithm for finding a sparse cut if λ_k is a constant for some constant k . This provides some theoretical justification to its empirical performance in image segmentation and clustering problems. The analysis can be extended to other graph partitioning problems, including multi-way partition, balanced separator, and maximum cut.

Joint work with Tsz Chiu Kwok, Yin Tat Lee, Shayan Oveis Gharan, Luca Trevisan.

Speaker: Jelani Nelson

Title: Dimensionality reduction via sparse matrices

Abstract: This talk will discuss sparse Johnson-Lindenstrauss transforms, i.e. sparse linear maps into much lower dimension which preserve the Euclidean geometry of a set of vectors. Both upper and lower bounds will be discussed, as well as applications to certain domains such as numerical linear algebra.

Speaker: Lorenzo Orecchia

Title: A Simple, Combinatorial Algorithm for Solving SDD Systems in Nearly-Linear Time

Abstract: Fast algorithms for Laplacian linear systems have found broad applications across both the theory and practice of computer science, playing a foundational role in scientific computing, machine learning, random processes, image processing, network analysis, and computational biology, among others. More recently, they have emerged as a powerful tool in the design of graph algorithms, enabling researchers to break longstanding algorithmic barriers for a wide range of fundamental graph problems, including finding maximum flows and multicommodity flows, generating random spanning trees, sparsifying graphs, routing in distributed systems, and finding sparse cuts and balanced separators.

Finding fast solvers for these systems has been an active area of research, culminating in algorithms that solve them in nearly-linear time. These algorithms are quite complex, as they rely on an intricate recursive construction of preconditioners, whose development required multiple fundamental innovations in spectral and combinatorial graph theory, graph algorithms, and computational linear algebra.

In this talk, I will present a very simple combinatorial algorithm that solves Laplacian systems in nearly-linear time. It uses very little of the machinery that previously appeared to be necessary for a such an algorithm. It does not require recursive preconditioning, or even the Chebyshev Method or Conjugate Gradient. After constructing a "nice" spanning tree of a graph associated to the linear system, the entire algorithm consists of the repeated application of a simple (non-recursive) update rule, which it implements using a lightweight data structure. The algorithm is numerically stable and can be implemented without the increased bit-precision required by previous solvers. As such, the algorithm presented has the fastest known running time under the standard unit-cost RAM model.

While the original paper explains the algorithm in terms of electrical flow, I will also present an optimization interpretation of the algorithm based on gradient descent and discuss its relationship with older combinatorial Laplacian solver and the Kaczmarz iterative solver.

This is joint work with Jonathan Kelner, Aaron Sidford and Zeyuan Zhu.

Speaker: Lorenzo Orecchia

Title: An Almost-Linear-Time Algorithm for Approximate Max Flow in Undirected Graphs, and its Multicommodity Generalizations

Abstract: In this talk, I will describe a new framework for approximately solving flow problems in capacitated, undirected graphs, and I will apply it to provide asymptotically faster algorithms for the maximum s-t flow and maximum concurrent multicommodity flow problems. For graphs with n vertices

and m edges, it allows us to find an ϵ -approximate maximum s - t flow in time $O(m^{1+o(1)} \epsilon^{-2})$, improving on the previous best bound of $mn^{1/3} \text{poly}(1/\epsilon)$. Applying the same framework in the multicommodity setting solves a maximum concurrent multicommodity flow problem with k commodities in $O(m^{1+o(1)} \epsilon^{-2} k^2)$ time, improving on the existing bound of $m^{4/3} \text{poly}(k, 1/\epsilon)$.

In describing these algorithms, I will discuss several new technical tools that may be of independent interest:

- * a non-Euclidean generalization of gradient descent, bounds on its performance, and a way to use this to reduce approximate maximum flow and maximum concurrent flow to oblivious routing;
- * the definition and efficient construction of a new type of flow sparsifier that ameliorates the longstanding gap between the algorithmic efficacy of sparsification on flow and cut problems; and
- * the first almost-linear-time construction of an $O(m^{o(1)})$ -competitive oblivious routing scheme.

This is joint work with Jonathan Kelner, Aaron Sidford, and Yin Tat Lee.

Speaker: Thomas Rothvoss

Title: Better Bin Packing Approximations using Discrepancy Theory

Abstract: For bin packing, the input consists of n items with sizes s_1, \dots, s_n in $[0, 1]$ which have to be assigned to a minimum number of bins of size 1. The seminal Karmarkar-Karp algorithm from '82 produces a solution with at most $\text{OPT} + O(\log^2 \text{OPT})$ bins.

We provide the first improvement in now 3 decades and show that one can find a solution of cost $\text{OPT} + O(\log \text{OPT} * \log \log \text{OPT})$ in polynomial time. This is achieved by rounding a fractional solution to the Gilmore-Gomory LP relaxation using the Entropy Method from discrepancy theory. The result is constructive via algorithms of Bansal and Lovett-Meka.

Speaker: Thomas Rothvoss (MIT)

Title: The matching polytope has exponential extension complexity

Abstract: A popular method in combinatorial optimization is to express polytopes P , which may potentially have exponentially many facets, as solutions of linear programs that use few extra variables to reduce the number of constraints down to a polynomial.

After two decades of standstill, recent years have brought amazing progress in showing lower bounds for the so called extension complexity.

However, the central question in this field remained wide open:

can the perfect matching polytope be written as an LP with polynomially many constraints?

We answer this question negatively. In fact, the extension complexity of the perfect matching polytope in a complete n -node graph is $2^{\Omega(n)}$.

Speaker: David Steurer

Title: Sum-of-Squares and the Unique Games Conjecture

Abstract: I will survey recent developments about the Unique Games Conjecture and sum-of-squares methods, in particular, connections to graph partitioning, polynomial optimization, proof complexity, and quantum information theory.

Speaker: David Steurer

Title: Analytical Approach to Parallel Repetition

Abstract: We propose an analytical framework for studying parallel repetition, a basic product operation for one-round two-player games. In this framework, we consider a relaxation of the value of a game, val_+ , and prove that for projection games, it is both multiplicative (under parallel repetition) and a good approximation for the true value, which implies parallel repetition bounds.

Using this framework, we can also give a short proof for the NP-hardness of Label-Cover $(1, \delta)$ for all $\delta > 0$, starting from the basic PCP theorem.

Concrete results of this work:

- A parallel repetition bound for projection games with value close to $\frac{1}{2}$. Previously, it was not known whether parallel repetition decreases the value of such games. This result implies tight inapproximability bounds for SET-COVER.
- An improved bound for few parallel repetitions of projection games, showing that Raz's counterexample is tight even for a small number of repetitions.

Joint work with Irit Dinur.

Speaker: Nisheeth Vishnoi

Title: An Interior Point Method for Directed Max-Flow

Abstract: I will talk about a recent result of Madry which shows how one can combine interior point methods for solving linear programs with Laplacian solvers to come up with new (and sometimes faster) algorithms for Directed Max-Flow.

Abstracts of FSTTCS 2013 Invited Talks

Speaker: Venkatesan Guruswami

Title: Polar codes: Reliable communication with complexity polynomial in the gap to Shannon capacity

Abstract: Shannon's monumental 1948 work laid the foundations for the rich fields of information and coding theory. The quest for *efficient* coding schemes to approach Shannon capacity has occupied researchers ever since, with spectacular progress enabling the widespread use of error-correcting codes in practice. Yet the theoretical problem of approaching capacity arbitrarily closely with polynomial complexity remained open except in the special case of erasure channels.

In 2008, Arikan proposed an insightful new method for constructing capacity-achieving codes based on channel polarization. In this talk, I will begin with a self-contained survey of Arikan's celebrated construction of polar codes, and then discuss our recent proof (with Patrick Xia) that, for all binary-input symmetric memoryless channels, polar codes enable reliable communication at rates within $\epsilon > 0$ of the Shannon capacity with block length (delay), construction complexity, and decoding complexity all bounded by a *polynomial* in the gap to capacity, i.e., by $\text{poly}(1/\epsilon)$. Polar coding gives the ***first explicit construction*** with rigorous proofs of all these properties; previous constructions were not known to achieve capacity with less than $\exp(1/\epsilon)$ decoding complexity.

We establish the capacity-achieving property of polar codes via a direct analysis of the underlying martingale of conditional entropies, without relying on the martingale convergence theorem. This step gives rough polarization (noise levels $\sim \epsilon$ for the "good channels"), which can then be adequately amplified by tracking the decay of the channel Bhattacharyya parameters. Our effective bounds imply that polar codes can have block length bounded by $\text{poly}(1/\epsilon)$. We also show that the generator matrix of such polar codes can be constructed in polynomial time by algorithmically computing an adequate approximation of the polarization process.

Speaker: Subhash Khot

Title: On Approximation Resistance of Predicates

Abstract: Constraint satisfaction problems are some of the most well-studied NP-hard problems, 3SAT being a prominent example. It is known by Hastad's 1997 result that 3SAT is "approximation resistant" in the following sense: given a near-satisfiable instance, a trivial algorithm that assigns random boolean values to the variables satisfies $7/8$ fraction of the constraints and no efficient algorithm can do strictly better unless $P = NP$!

3SAT is a CSP that corresponds to the ternary OR predicate. In general, a CSP has constraints given by some fixed predicate $P: \{0,1\}^k \rightarrow \{\text{True}, \text{False}\}$ (on possibly negated variables) and the predicate is called approximation resistant if, on a near-satisfiable instance, it is computationally hard to perform strictly better than a random assignment.

The quest to understand approximation resistance has played a central role in the theory of probabilistically checkable proofs (PCPs) and hardness of approximation. This talk will give a survey of the topic, including recent work towards a complete characterization of approximation resistance (i.e. a necessary and sufficient condition on the predicate that makes the corresponding CSP approximation resistant).

