ACO Colloquium Series

Wednesday, February 6, 2008

Gary Miller
Carnegie-Mellon University

Image Segmentation using Spectral Graph Theory

4:30pm, Skiles 255

In this talk we present a new image segmentation algorithm, Spectral Rounding (SR), and a fast solver. The combination is used for segmenting 2D and 3D images. Applying SR to the Berkeley data base of human segmented images, and medical examples such as tumors in mammograms and 3D retinal scans gives robust quality segmentations. When used with our fast solver SR is nearly nearly time.

The key idea in SR is to view an image as a 2D mattress of springs. Two neighboring pixels are connected by a spring where the spring constant is determined by local similarity in the pixel intensity. Shi and Malik proposed the important idea of using the fundamental modes of vibration of this mattress, the eigenvectors, to segment the image. The straightforward method for partitioning a graph using its eigenvectors, however, does not seem to work well in practice.

We propose a relaxation method based on eigenvectors for finding these graph cuts. At each round a few fundamental eigenvectors are computed, from which the spring constants are updated and these eigenvectors are recomputed using the new spring constants. Thus the spring constants are successively readjusted until the mattress disconnects, an image segmentation.

SR compares favorably with hand-segmented images from the Berkeley database and the normalized cut metric. We also show convergence in general and termination for several important cases.

The second issue addressed is fast algorithms for finding the associated eigenvectors and solving related linear systems. This is a critical issue because modern 3D medical images may contain a billion nodes (voxels). A related and important first step to finding eigenvector and of interest on its own is solving 2D and 3D Laplacians. For instance, Siemens uses Laplacians for their new assisted image segmentation algorithm. We present the first linear-time algorithm for 2D and more general planar Laplacians.

This represents joint work with Yiannis Koutis and David Tolliver.

Monday, November 12, 2007

Balaji Prabhakar
Stanford University

Turbo-Counters: Efficient network traffic measurement using a sparse graph counter architecture

3:30pm, Klaus 1116

Measuring network flow sizes is important for accounting/billing, network forensics and security. Per-flow measurement is considered hard because it requires that many counters be updated at a very high speed, and a flow-to-counter perfect hash function. Therefore, current approaches aim to obtain approximate flow counts; that is, to detect large "elephant" flows and then measure their sizes. We present a novel method for per-flow traffic measurement that is fast, highly memory efficient and accurate. At the core of this method is an architecture, which we call "counter braids." We present two results: (i) The optimality of counter braids, in the sense that the total number of bits needed store flow sizes losslessly equals the entropy lower bound. Since network traffic has an entropy rate less than 3 bits per flow, this makes possible an implementation in an on-chip SRAM. (ii) A low-complexity message-passing estimation algorithm, that recovers flow sizes with vanishing error at link speed. Evaluation on Internet traces demonstrates that almost all flow sizes are estimated error-free with only 5 bits per flow. Joint work with: Sarang Dharmapurikar, Abdul Kabbani, Yi Lu, Andrea Montanari.

Thursday, October 18, 2007

Sandy Irani
Computer Science Department, UC Irvine

The power of quantum systems on a line

4:30pm, Skiles 269

In this talk, we discuss the computational strength of finite dimensional quantum particles arranged on a line. We prove that adiabatic computation is equivalent to standard quantum computation even when the adiabatic quantum system is restricted to 2-local interactions of nearest neighbors on a line. The particles in this construction require 9 states per particle. We then adapt this construction to show that the 2-local Hamiltonian for 12 state particles on a line is QMA-complete. QMA is the quantum analog of NP. This result contrasts with the classical analog in which one-dimensional Max-2-SAT with nearest neighbor constraints is in known to be in P. Similar results were obtained independently by Aharonov, Gottesman and Kempe. The work appears jointly in FOCS 2007.
Refreshments at 4:00PM in Skiles 236

Thursday, October 16, 2007

Leonid Gurvits
Los Alamos National Labs

Van der Waerden/Schrijver-Valiant like Conjectures and Stable Homogeneous

4:30pm, Skiles 255

Refreshments at 4:00PM in Skiles 236

Monday, October 1, 2007

Mohit Singh
Tepper School of Business, Carneigie Mellon University

Iterative Methods in Combinatorial Optimization

4:30pm, Klaus 1116E

Linear programming has been a successful tool in combinatorial optimization to achieve polynomial time algorithms for problems in P and also to achieve good approximation algorithms for problems which are NP-hard. We demonstrate that iterative methods give a general framework to analyze linear programming formulations of combinatorial optimization problems. We show that iterative methods are well-suited for problems in P and lead to new proofs of integrality of linear programming formulations for various problems in P. This understanding provides us the basic groundwork to address various problems that are NP-hard and to achieve good approximation algorithms. In this talk, we focus on degree bounded network design problems. The most well-studied problem in this class is the Minimum Bounded Degree Spanning Tree problem. We present a polynomial time algorithm that returns a spanning tree of optimal cost such that the degree of any vertex in the tree exceeds its degree bound by at most an additive one. This generalizes a result of Furer and Raghavachari to weighted graphs, and thus settles a 15-year-old conjecture of Goemans affirmatively. This is essentially the best possible result for this problem. For degree constrained versions of more general network design problems, we obtain strong bi-criteria approximation algorithms using the iterative method.

Tuesday, March 6, 2007

Jeff Kahn
Mathematics, Rutgers University

Correlation Inequalities

4:30pm, Skiles 255

Correlation inequalities --- e.g. the Harris, FKG and BK inequalities --- are basic tools in probability and have also had some

beautiful combinatorial consequences. Here we will mention a few relatively recent inequalities, but mainly want to give some indication of how little we actually know about such things.
Refreshments at 4:00PM in Skiles 236

Tuesday, February 13, 2007

Yuri Nesterov
Université catholique de Louvain

Controllable random permutations

4:30pm, Skiles 255

In this talk, we propose a new procedure for generating random permutations. We discuss applications of this technique to some scheduling problems. In particular, we show that by applying this procedure to a single machine, it is possible to ensure a desired average completion time for every job, or to prove that this is not feasible. The parameters of the procedure can be found in polynomial time by solving a simple nonlinear convex optimization problem.
Refreshments at 4:00PM in Skiles 236

Tuesday, December 5th, 2006

Sergei Izmalkov
Department of Economics, MIT

Perfect Implementation of Normal-Form Mechanisms

4:30pm, Skiles 255

Privacy and trust affect our strategic thinking, yet they have not been precisely modeled in mechanism design. In settings of incomplete information, traditional implementations of a normal-form mechanism ---by disregarding the players' privacy, or assuming trust in a mediator--- may not be realistic and fail to reach the mechanism's objectives. We thus investigate implementations of a new type. We put forward the notion of a perfect implementation of a normal-form mechanism M: in essence, an extensive-form mechanism exactly preserving all strategic properties of M, without relying on a trusted party or violating the privacy of the players. We prove that any normal-form mechanism can be perfectly implemented via envelopes and an envelope-randomizing device (i.e., the same tools used for running fair lotteries or tallying secret votes).

Tuesday, November 7th, 2006

Bertrand Guenin
University of Waterloo

Towards a proof of Seymour's 1-flowing conjecture

4:30pm, Skiles 269

The well known max-flow min-cut theorem states that the maximum amount of flow that can be sent from vertex s to a vertex t in a graph (with capacities) is equal to the capacity of smallest cut which separates s and t. While the notion of flows in graphs extends naturally to binary matroids, the max-flow min-cut relation does not hold for binary matroids in general. However, Seymour conjectured that the aforementioned minimax relation holds as long as the binary matroids do not contain any one of three special obstructions. This conjecture if true would generalize many classical results on multi-commodity flows and matchings. Our approach is to try to give a structural characterization of the binary matroids for which the minimax relation holds. I will review known cases of the conjecture and give a brief sketch of our strategy for solving the conjecture. A more technical description of the tools we are developing will be presented on November 10 during the Combinatorics seminar. This is joint work with Irene Pivotto and Paul Wollan.
Refreshments at 4:00PM in Skiles 236

Friday, May 4th, 2006

Jeff Kahn
Mathematics, Rutgers University

Many Hamiltonian Cycles

11:05am, Skiles 255

We'll begin with the following theorem, which proves a conjecture of Sarkozy, Selkow and Szemeredi, and try to use it as an excuse to talk about other things (e.g., Bregman's Theorem, entropy, the "incremental random method," statistics physics ...).
Theorem Any n-vertex Dirac graph (i.e., graph of minimum degree at least n/2) contains at least (2-o(1))^{-n}n! Hamiltonian cycles.
Joint with Bill Cuckler.

Friday, September 23rd, 2005

Benny Sudakov
Mathematics, Princeton University

Additive Approximation for Edge-Deletion Problems

4:05, Skiles 255

A graph property is monotone if it is closed under removal of vertices and edges. In this talk we consider the following edge-deletion problem; given a monotone property P and a graph G, compute the smallest number of edge deletions that are needed in order to turn G into a graph satisfying P. We denote this quantity by E_P(G). Our first result states that for any monotone graph property P, any \epsilon> 0 and n-vertex input graph G one can approximate E_P(G) up to an additive error of \epsilon n^2. Given the above, a natural question is for which monotone properties can one obtain better additive approximations of E_P(G). Our second main result essentially resolves this problem by giving a precise characterization of the monotone graph properties for which such approximations exist. We will show that for any dense monotone property, that is, a property for which there are graphs on n vertices with \Omega (n^2) edges that satisfy it, it is NP-hard to approximate E_P(G) up to an additive error of n^{2-\delta}, for any fixed positive \delta. The proof requires several new ideas and involves tools from Extremal Graph Theory together with spectral techniques. Interestingly, prior to this work it was not even known that computing E_P(G) precisely for dense monotone properties is NP-hard. We thus answer (in a strong form) a question of Yannakakis raised in 1981. (Joint work with N. Alon and A. Shapira.)

Wednesday, September 7th, 2005

Ashish Goel
Stanford University

Algorithmic Self-Assembly: Models and Problems

11:05, Skiles 255

DNA Self-assembly has emerged as an important technique for molecular computation and nano-technology. At these scales, self- assembly is governed by simple (and local) probabilistic rules for growth, making it amenable to algorithmic techniques. We will discuss two important challenges in algorithmic self-assembly: robustness and efficiency. This talk will present recent results, and also attempt to provide a road-map of open problems.

Friday, October 29th, 2004

Belá Bollobás
University of Memphis, TN and Fellow of Trinity College, Cambridge, UK

Voronoi Percolation

4:00pm, Skiles 269

Friday, November 5, 2004

Alan Frieze
Carnegie Mellon University

One and two stage minimum spanning tree problems: average case analysis

4:00pm, Skiles 269

Let the edge weights of a graph G be given independent uniform [0,1] random costs. We review some old results about the expected value of the minimum spanning tree. We then consider a 2-stage problem where independent weights are given for Monday and Tuesday. On Monday we might buy some edges knowing only the distribution of the edge weights for Tuesday and then buy the remaining edges on Tuesday. We present some results on this and on a directed version.

Wednesday, February 11th, 2004

Yan Ding
College of Computing, Georgia Tech

Constant Round Oblivious Transfer in the Bounded Storage Model

4:30pm, Skiles 255

The bounded storage model, introduced by Maurer, is an alternative to the standard complexity based model for cryptography. In contrast to the complexity based model, the bounded storage model assumes that the adversary is space-bounded, and provides provable information-theoretic security by employing a public random string R whose length exceeds the space bound. Security is guaranteed against a malicious adversary who remembers almost all information about R, while an honest party is only required to store a small fraction of R. Oblivious transfer is a fundamental cryptographic primitive on which many two-party and multi-party cryptographic protocols can be based. It involves two parties, Alice who has two secrets s_0 and s_1, and Bob who has a secret bit b. Roughly speaking, in a secure oblivious transfer protocol, Alice sends s_0 and s_1 to Bob in such a way that (1) Bob receives s_b but learns nothing about the other secret, and (2) Alice learns nothing about b. We construct a 5-round protocol for oblivious transfer in the bounded storage model. This improves on previous works that required polynomially many rounds. Our protocol also significantly improves upon the previous ones in terms of total bit communication complexity. The main ingredients of our construction are powerful tools from pseudo randomness, an area that is concerned with generating "random looking" objects with no or little randomness. In particular, our construction uses randomness extractors, averaging samplers, and almost t-wise independent permutations. This talk will be self-contained, and assume no previous knowledge in cryptography and pseudo randomness. This is joint work with Danny Harnik, Alon Rosen and Ronen Shaltiel.