ABSTRACTS

Online algorithm for a bandwidth utilization problem

Sanjeev Arora
Dept. of Computer Science
Princeton University

Karp, Papadimitriou, and Shenker recently introduced the following model that captures the task of a sender trying to send messages over a congested network. At time t the total available bandwidth is b_t, which is unknown to the sender except it knows that bandwidths at successive time periods satisfy a weak continuity relation: $b_t \in [b_{t-1}/\mu, \mu b_{t-1}]$, where μ is some constant. The sender elects to send x_t bits. If $x_t \leq b_t$ then all get delivered, and if $x_t > b_t$ then none get delivered. The goal is to maximize $U = \sum_t x_t$. Note that this quantity is upperbounded by $B = \sum_t b_t$. We call U/B the *performance ratio.*
Karp et al. showed that for every deterministic online algorithm there is a sequence of bandwidths $\{b_t\}$ such the performance ratio is at most $1/\mu$, and that there is a simple algorithm that achieves this ratio. They could not do a similar analysis for randomized algorithms.
We show a randomized online algorithm that achieves a performance ratio $O(1/\log \mu)$ and prove that no other algorithm can do better.
Our algorithm is a variant of a classic strategy called *multiplicative increase multiplicative decrease.* We discuss implications of this fact, including morals for designers of network protocols.
Joint work with William Brinkman.

Metric Geometry of Counting

Alexander Barvinok
Dept. of Mathematics
University of Michigan

We describe general methods to obtain fast (polynomial time) estimates of the cardinality of a combinatorially defined set via solving some randomly generated optimization problems on the set. Examples include enumeration of perfect matchings in graps, bases in matroids, forests, spanning subgraphs, etc. Geometrically, we estimate the cardinality of a subset of the Boolean cube via the average distance from a point in the cube to the subset.
Joint work with A. Samorodnitsky.

The natural Workstealing Algorithm is stable

Petra Berenbrink
Dept. of Computer Science
University of Warwick

In this paper we analyse a very simple dynamic work-stealing algorithm. In the work-generation model, there are n *generators* which are *arbitrarily* distributed among a set of n *processors*. The distribution of generators is arbitrary — generators may even move at the beginning of each time step. During each time-step, each generator may generate a unit-time task which it inserts into the queue of its host processor. It generates such a task independently with probability λ. After the new tasks are generated, each processor removes one task from its queue and services it. Clearly, the work-generation model allows the load to grow more and more imbalanced, so, even when $\lambda < 1$, the system load is unbounded. The natural work-stealing algorithm that we analyse is widely used in practical applications and works as follows. During each time step, each *empty* processor (with no work to do) sends a request to a randomly selected other processor. Any *non-empty* processor having received at least one such request in turn decides (again randomly) in favour of one of the requests. The number of tasks which are transferred from the non-empty processor to the empty one is determined by the so-called *work-stealing function* f. In particular, if a processor that accepts a request has ℓ tasks stored in its queue, then $f(\ell)$ tasks are transferred to the currently empty one. A popular work-stealing function is $f(\ell) = \lfloor \ell/2 \rfloor$, which transfers (roughly) half of the tasks. We analyse the *long-term behaviour* of the system as a function of λ and f. We show that the system is *stable* for any constant generation rate $\lambda < 1$ and for a wide class of functions f. Most intuitively sensible functions are included in this class (for example, every function $f(\ell)$ which is $\omega(1)$ as a function of ℓ is included). We give a quantitative description of the functions f which lead to stable systems. Furthermore, we give *upper bounds* on the average system load (as a function of f and n). Our proof techniques combine Lyapunov function arguments with domination arguments, which are needed to cope with dependency.

Approximation Hardness of Bounded Degree MIN-CSP and MIN-BISECTION

Piotr Berman
Dept. of Computer Science
Universerstiy of Bonn

We consider bounded occurrence (degree) instances of a minimum constraint satisfaction problem MIN-LIN2 and a MIN-BISECTION problem for graphs. MIN-LIN2 is an optimization problem for a given system of linear equations mod 2 to construct a solution that satisfies the minimum number of them. E3-OCC-MIN-E3-LIN2 is the bounded occurrence (degree) problem restricted as follows: each equation has exactly 3 variables and each variable occurs in exactly 3 equations. Clearly, MIN-LIN2 is equivalent to another well known problem, the Nearest Codeword problem, and E3-OCC-MIN-E3-LIN2 to its bounded occurrence version. MIN-BISECTION is a problem of finding a minimum bisection of a graph, while 3-MIN-BISECTION is the MIN-BISECTION problem restricted to 3-regular graphs only. We show that, somewhat surprisingly, these two restricted problems are exactly as hard to approximate as their general versions. In particular, an approximation ratio lower bound for E3-OCC-MIN-E3-LIN2 (bounded 3-occurrence 3-ary Nearest Codeword problem) is equal to MIN-LIN2 (Nearest Codeword problem) lower bound $n^{\Omega(1)/\log\log n}$. Moreover, an existence of a constant factor approximation ratio (or a PTAS) for 3-MIN-BISECTION entails existence of a constant approximation ratio (or a PTAS) for the general MIN-BISECTION.
Joint work with Marek Karpinski.

Slow Mixing for H-Colorings of the Hypercubic Lattice

Christian Borgs
Microsoft Research
Redmond

An H coloring of a simple graph G is map from G to H that maps each edge in G into an edge in H. It is known that the problem of deciding whether such an H-coloring exists is NP-complete if H has no loops and is not bipartite (Hell and Nešetřil, 1990), and polynomial otherwise. The counting problem, i.e. the problem of counting the number of H-colorings of a graph G, is $\sharp$P-complete if H is neither the completely looped complete graph, K_n^{loop}, nor the complete bipartite graph, $K_{n,m}$, and polynomial otherwise (Dyer and Greenhill, 2000). Motivated by this result, we call H *trivial if* $H = K_n^{\text{loop}}$ *or* $H = K_{n,m}$.
In this work, we study random H-colorings of rectangular subsets of the hypercubic lattice $\mathbf{Z}^d$, with weight $\lambda_i \in (0, \infty)$ for the color i. We consider quasi-local Markov chains on a periodic box of even side length L, that is, Markov chains that do not change more than a fraction $\rho < 1$ of the sites in the box in any single move. For any finite, connected, non-trivial H, we show that there are weights $\{\lambda_i\}$ such that all quasi-local reversible ergodic Markov chains have slow mixing in the sense that the mixing time is exponential in $L^{d-1}/(\log L)^2$. Under the same conditions, we prove phase coexistence in the sense that there are at least two extremal Gibbs states. We also prove that, for a large subclass of graphs H, one can choose weights $\{\lambda_i\}$ such the corresponding Gibbs measure has exponentially fast spatial mixing.
Joint work with Jennifer T. Chayes, Martin Dyer, and Prasad Tetali.

Connectivity among H-colourings of graphs

Graham Brightwell
Dept. of Mathematics
London School of Economiecs

An *H-colouring* of a graph G is a homomorphism from G to H; and $\hom(G, H)$ denotes the set of all H-colourings of G. Two H-colourings are deemed to be adjacent if they differ on only one vertex of G; we are interested in when $\hom(G, H)$ is connected: this is an obvious necessary condition for single-site Glauber dynamics to be rapidly mixing for $\hom(G, H)$.
Jerrum had observed that, in the special case where H is the complete graph K_n, $\hom(G, H)$ is connected for all graphs G of maximum degree at most $n-2$, but not for all graphs of maximum degree $n-1$. Generally we say that H is *d-mobile* if $\hom(G, H)$ is connected for all G of maximum degree at most $d-2$: so K_d is d-mobile but not $(d+1)$-mobile. We conjecture that no d-colourable graph is $(d+1)$-mobile. We prove this in the case $d = 3$, and also prove the weaker result that no d-colourable graph is $(2d-1)$-mobile. Our proof for $d = 3$ uses the notion of the *circular chromatic number* of H; for larger d we use a generalisation of this concept to higher dimensions.
Joint work with Peter Winkler.

The Phase Transition in the Random Partition Problem

Jennifer Chayes
Microsoft Research
Redmond

The integer partition problem is a canonical NP-complete problem of combinatorial optimization. We show that the random version of this problem has a phase transition and establish the behavior of the model near the transition. In particular, we show that the phase transition is discontinuous or "first-order," in contrast to the phase transitions established in other combinatorial models such as the random graph and the 2-satisfiability problem. We also discuss recent suggestions that the order of the phase transition may be related to the hardness of the problem.
Joint work with C. Borgs and B. Pittel.

Random Garphs which model the internet

Colin Cooper
Dept. of Mathematical & Computing Sciences
University of London
Goldsmiths College

We consider the degree sequence of a general model of web graphs. For a wide range of the parameters of the model, the degree sequence obeys a power law whose parameter is a function of these parameters.

On Certain Property Testing Algorithms

Artur Czumaj
Dept. of Computer and Information Science
New Jersey Institute of Technology

We introduce a new framework for analyzing property testing algorithms. Informally, our framework can be applied to decision problems that can be described as a pair of "bases" and "constraints," and the instance is accepted if there is a basis which is not "violated" by any constraint. We show, again informally, that if for a given problem it is possible to define the bases to be of small size, then the problem possesses a constant-time testing algorithm. We present our approach in a rather generic framework that has simple formulation and can be applied to a large variety of problems.
We apply our framework to obtain property testing algorithms for the most representative and most widely studied problems of graph coloring, clustering, some algebraic problems, some problems related to linear and mathematical programming, and for some covering problems.
Our approach, besides its generality and simplicity, leads in many cases to either new or improved results.
Joint work with Christian Sohler.

Approximation Hardness of TSP with Bounded Metrics

Lars Engebretsen
Laboratory for Computer Science
MIT

The general asymmetric TSP with triangle inequality is known to be approximable only to within an $O(\log n)$ factor, and is also known to be approximable within a constant factor as soon as the metric is bounded by a constant. In this talk, we discuss techniques for proving lower bounds on the approximability of TSP with bounded metrics. In particular, we first give lower bounds for the asymmetric and symmetric versions of TSP with distances one and two by the means of a gadget reduction from a problem called *Hybrid*, consisting of a system of linear equations mod 2 with either two or three variables per equation and exactly three occurrences of each variable. We also note that the construction used by Papadimitriou and Vempala to prove their recently announced lower bounds on the approximability of the general TSP with triangle inequality can be modified slightly to give comparable lower bounds also for the case when the metric is bounded by a small constant.
Joint work with Marek Karpinski.

Sampling k-Uniform Hypergraphs and Design of PTASs for Dense Instances of Min-CSP

W. Fernandez de la Vega
Laboratoire de Recherche en Informatique
Université Paris-Sud

We introduce a new sampler technique for k-uniform hypergraphs and apply it to design the first *polynomial time approximation schemes* (PTASs) for dense instances of MIN-Ek-LIN2 (the problem of minimising the number of satisfied equations within a system of linear equations mod 2 with exactly k variables per equation) and dense instances of MIN-Ek-SAT.
Joint work with C. Bazgan and M. Karpinski.

Arc-Disjoint Paths in Expander Digraphs

Alan Frieze
Mathematical Sciences Dept.
Carnegie Mellon University

Given a digraph $D = (V, A)$ and a set of κ pairs of vertices in V, we are interested in finding for each pair (x_i, y_i), a directed path connecting x_i to y_i, such that the set of κ paths so found is arc-disjoint. For arbitrary graphs the problem is $\mathcal{NP}$-complete, even for $\kappa = 2$.
We present a polynomial time randomized algorithm for finding arc-disjoint paths in an r-regular expander digraph D. We show that if D has sufficiently strong expansion properties and r is sufficiently large then *all* sets of $\kappa = \Omega(n/\log n)$ pairs of vertices can be joined. This is within a constant factor of best possible.
Joint work with Tom Bohman.

Connectedness of the bounded-degree star process

Catherine Greenhill
Dept. of Mathematics & Statistics
University of Melbourne

A graph process starts with an empty graph and at each step adds an edge or edges, chosen according to some probabilistic rule. For fixed d, the star d-process chooses a vertex v of minimum degree i, uniformly at random, and then chooses d-i vertices of degree less than d, uniformly at random, and joins each of these to v. Rucinski and Wormald proved that the resulting graph is asymptotically almost surely d-regular (when dn is even). We prove that the final graph is asymptotically almost surely connected for d at least 3, and is a.a.s. d-connected for large enough d (d at least 15 should do).
Joint work with Andrzej Ruciński and Nicholas C. Wormald.

Polynomial-time Approximation Algorithms for Preemptive Resource Constrained Scheduling and Fractional Graph Coloring.

Klaus Jansen
Dept. of Computer Science
Universität Kiel

We study resource constrained scheduling problems where the objective is to compute feasible preemptive schedules minimizing the makespan and using no more resources than what are available. We present approximation algorithms along with some inapproximibility results showing how the approximability of the problem changes in terms of the number of resources. All the results are based on linear programming formulations (though with exponentially many variables) that are called fractional covering problems. Furthermore we show some interesting connections between resource constrained scheduling and (multi - dimensional, multiple-choice, and cardinality constrained) variants of the classical knapsack problem. Finally we present applications of the above results in fractional graph coloring and multiprocessor task scheduling.
Joint work with Lorant Porkolab, Imperial College London.

Theoretical Analysis of Evolutionary Algorithms

Thomas Jansen
FB Informatik II
Universität Dortmund

Evolutionary algorithms are randomized search heuristics that are often used for optimization of pseudo-boolean functions $f\colon \{0,1\}^n \rightarrow \mathbf{R}$. They are well-established in practice and intensively empirically investigated since the 1980s. However, their theoretical foundation is still unsatisfying. This is especially true for evolutionary algorithms that use crossover. Here, three examples are presented where one can prove that appropriate genetic algorithms with crossover out-perform by far mutation-based evolutionary algorithms. The first example is diversity oriented and proves a small polynomial expected running time for a steady-state GA with uniform crossover whereas mutation-based EAs have super-polynomial expected running time. The second example even proves an exponential gap between a GA with 1-point crossover and mutation-based EAs. Finally, a third example proves the same for a GA with uniform-crossover. The examples are based upon artificial example functions that are all well-structured, understandable and provide some insight. They are considered to be helpful first steps towards a rigorous analysis of evolutionary algorithms on natural problems.
Joint work with Ingo Wegener.

Approximating the Permanent (part I)

Mark Jerrum
Dept. of Computer Science
University of Edinburgh

This two-part presentation (with Vigoda) develops a fully-polynomial randomized approximation scheme for computing the permanent of an arbitrary matrix with non-negative entries. Part I sets the scene by reviewing an existing MCMC approach to approximating the permanent, proposed by Broder and made rigorous by Jerrum and Sinclair using the "canonical paths" argument. The limitations of the existing approach are described, and an obstacle to further progress identified.
Joint work with Alistair Sinclair.

Approximating the Permanent (part II)

Eric Vigoda
Dept. of Computer Science
University of Edinburgh

This two-part presentation (with Jerrum) develops a fully-polynomial randomized approximation scheme for computing the permanent of an arbitrary matrix with non-negative entries. Part II describes how to modify the existing Markov chain—by applying carefully chosen weights to configurations of unmatched vertices—in order to achieve rapid mixing for all problem instances. The weights may be approximated by an iterative procedure in which MCMC is used to adjust the weights at each step.
Joint work with Alistair Sinclair.

What Can you do in one or two passes?

Ravi Kannan
Dept. of Computer Science
Yale University

There are many applications in which the input data is too large to be stored in RAM. In such cases, it makes sense to restrict the number of passes one is allowed to make through the entire data because a pass which has to be from disk is costly. We study problems which can be approximately solved by making one or two passes through the data in which we sample a small part of the data and then compute on the sample in time polynomial in only the size of the sample. In "blind sampling", one samples (usually uniformly at random) without first reading. In other algorithms, the sampling probabilities are based on one read of the data. This gives us considerable advantage in many problems like the max cut and other discrete problems as well as Principal Component analysis and some Information Retreival problems.

Distributed $O(\Delta \log n)$ -edge-coloring algorithm

Michal Karoński
Dept. of Mathematics
University of Poznan

We consider a problem of edge-coloring of a graph in a distributed model of computations. In our model a network is represented by an undirected graph $G = (V, E)$ where each vertex represents a processor of the network and an edge corresponds to a connections between processors. We assume full synchronization of the network: in every step, each processor sends messages to all its neighbors, receives messages from all of its neighbors, and can perform some local computations. However, we insist that the local computations must be performed in time which is polynomial in the size of the graph. By default, all processors have different IDs, each proccesor knows $|V|$, the number of vertices in G, and $\Delta(G)$, the maximal degree in G. In the edge-coloring problem the goal of a distributed algorithm is to properly color the edges of G in a polylogarithmic (in $n = |V|$) number of steps. In our talk, we present a distributed algorithm which colors edges of graph in $O(\Delta \log n)$ colors. Our approach is based on computing a family of spanners of G. It turns out that this family can be used to color a constant fraction of edges of G using $O(\Delta)$ colors. Iterating this process $O(\log n)$ steps leads to a proper coloring of E. However in each iteration a palette of $O(\Delta)$ new colors is needed. Spanners were previously successfully used by Hańćkowiak, Karoński and Panconesi, to design a distributed algorithm for a maximal matching problem.
Joint work with A. Czygrinow, M. Hańćkowiak.

Approximability of Dense Nearest Codeword Problem

Marek Karpinski
Dept. of Computer Science
University of Bonn

We design a *polynomial time approximation scheme* (PTAS) for the dense instances of *Nearest Codeword Problem* (NCP). The problem can be formulated as a *linear feasibility* problem of constructing an assignment x for a given system of linear equations mod 2, which minimizes the number of unsatisfied equations. The Dense NCP was known to be NP-hard in an exact setting. The general problem is known to have exceedingly high lower approximation bound of $n^{\Omega(1)/loglogn}$ (Dinur, Kindler, Raz, Safra, 2000), and an existence of a PTAS on dense instances comes as a surprise. The technique of solution depends on a method of approximating Smooth Polynomial Integer Programs (Arora, Karger and Karpinski, 1995), and a new density sampler technique for graphs and k-uniform hypergraphs developed recently by Bazgan, Fernandez de la Vega and Karpinski, 2000. Despite an importance of the general NCP problem, and its many motivations, not much was known about "good" approximation ratio algorithms, better than of order n, and this for arbitrary fields. Only recently the first polynomial time algorithm with sublinear approximation ratio $O(n/logn)$ was designed for the general problem by Berman and Karpinski, 2001. A challenging problem remains to design a better approximation algorithm which works on general instances of NCP. $\square$

Coalescing Particles on a Tree

Claire Kenyon
Laboratoire de Recherche en Informatique
Université Paris-Sud

The following problem is related to the average-case analysis of distributed updates on trees. Consider a perfect binary tree of height h. At time 0, we begin with a particle at each tree node. At each positive integer time, one of the remaining particles is chosen at random and moved up to its parent node, coalescing with any particle that might already be there. How long does it take until all particles coalesce (at the root)?
Joint work with Alistair Sinclair.

Algorithms for Minimizing Preemptive Weighted Flow Time

Sanjeev Khanna
Dept. of Computer and Information Science
University of Pennsylvania

We present the first approximation schemes for minimizing weighted flow time on a single machine with preemption. Our first result is an algorithm that computes a $(1+\epsilon)$-approximate solution for any instance of weighted flow time in $n^{O(\log W \log P/\epsilon^3)}$ time; here P is the ratio of maximum job processing time to minimum job processing time, and W is the ratio of maximum job weight to minimum job weight. This result directly gives a quasi-PTAS for weighted flow time when P and W are poly-bounded, and a PTAS when they are both bounded. We strengthen the former result to show that in order to get a quasi-PTAS it suffices to have just one of P and W to be poly-bounded. Our result provides a strong evidence that the weighted flow time problem has a PTAS. We note that the problem is strongly NP-hard even for bounded P and W. We next consider two important special cases of weighted flow time, namely, when P is bounded and W is unrestricted, and when the weight of a job is inverse of its processing time, refered to as the stretch metric. For both cases we obtain a PTAS by combining a novel partitioning scheme with our PTAS for the case of bounded P and W.
Joint work with Chandra Chekuri.

Approximating Minimum Size 2-Connectivity Problems using Local Search

Piotr Krysta
MPI für Informatik

We study the problem of finding the minimum size 2-edge-connected spanning subgraph. This problem is NP-hard (even on cubic planar graphs) and Max SNP-hard in general. We show that the minimum 2-edge-connected subgraph problem can be approximated to within $\frac{4}{3} - \epsilon$ for general graphs, improving upon the recent result of Vempala and Vetta (APPROX 2000). The significance of this result follows from its relations to the long standing $\frac{4}{3}$ metric TSP conjecture, due to Goemans (1995). Better approximations are obtained for planar graphs and for cubic graphs. We also consider some generalizations of the 2-edge-connected spanning subgraph problem. It is important to note that most of our algorithms use local search paradigm as the main method or as a subroutine. In the case of cubic graphs, our results imply a new upper bound on the integrality gap of the natural linear programming formulation for the 2-edge-connected spanning subgraph problem.
Joint work with A. Anil Kumar.

The RPR^2 rounding technique for semidefinite programs

Michael Langberg
Dept. of Computer Science and Applied Mathematics
Weizmann Institute of Science

Several combinatorial optimization problems can be approximated using algorithms based on semidefinite programming. In many of these algorithms a semidefinite relaxation of the underlying problem is solved yielding an optimal vector configuration $v_1 \dots v_n$. This vector configuration is then *rounded* into a $\{0,1\}$ solution. We present a procedure called RPR^2 (Random Projection followed by Randomized Rounding) for rounding the solution of such semidefinite programs. We show that the *random hyperplane* rounding technique introduced by Goemans and Williamson, and its variant that involves *outward rotation* are both special cases of RPR^2. We illustrate the use of RPR^2 by presenting two applications. For Max-Bisection we improve the approximation ratio. For Max-Cut, we improve the tradeoff curve (presented by Zwick) that relates the approximation ratio to the size of the maximum cut in a graph.
Joint work with Uriel Feige.

Routing random calls on graphs

Malwina Luczak
Mathematical Institute
University of Oxford

We are given a complete graph and a sequence of calls uniformly distributed over the edges. For each call $\{v, u\}$ in turn, the call is routed on the direct link if possible; and otherwise d nodes are selected uniformly at random from $V \setminus \{v, u\}$ and the call is routed via one of these nodes if possible. The *first fit dynamic alternative routing* algorithm FDAR chooses the first possible alternative route. The *balanced dynamic alternative routing* algorithm BDAR chooses an alternative route which minimises the maximum of the current loads on its two links. We compare the asymptotic blocking probability achieved by these algorithms. We further consider some extensions to non-complete graphs and asymmetric distributions of calls.

Decomposition, Swapping and Mean-Field Models

Dana Randall
School of Mathematics
Georgia Institute of Technology

Simulated tempering is a compelling Markov chain heuristic used for random sampling when other Markov chains are known to be slow. The idea is to enhance the state space with a parameter modeling temperature, and to allow the temperature to vary during the simulation. At high temperature bottlenecks (which cause slow mixing) disappear, mixing occurs, and lowering the temperature recovers the stationary distribution of interest. The swapping algorithm is a variant of this method. Recently Madras and Zheng analyzed the swapping algorithm on two bimodal distributions, including the mean-field Ising model, and showed that it is efficient. Their proof utilizes the decomposition method in novel ways. We extend these results to show that the swap algorithm is efficient for some asymmetric distributions as well.

Quantum Algorithms for Some Instances of the Hidden Subgroup Problem

Miklos Santha
Laboratoire de Recherche en Informatique
Université Paris-Sud

In the first part of the talk we give a survey on the status of the hidden subgroup problem, and in particular we sketch an efficient quantum algorithm for the Abelian case. In the second part we show that certain special cases of the non-Abelian case can also be solved in polynomial time by a quantum algorithm. These special cases involve finding hidden normal subgroups of solvable groups and permutation groups, finding hidden subgroups of groups with small commutator subgroup and of groups admitting an elementary Abelian normal 2-subgroup of small index or with cyclic factor group.

Joint work with G. Ivanyos and F. Magniez.

Judicious partitions of graphs and hypergraphs

Alex Scott
Dept. of Mathematics
University College London

Many classical partitioning problems ask for the maximum or minimum of a given quantity over partitions of a graph G. For instance, the classical Max Cut problem asks for the maximum of $e(V_1, V_2)$ over partitions $V(G) = V_1 \cup V_2$, or equivalently the minimum of $e(V_1) + e(V_2)$. *Judicious* partitioning problems ask for some quantity to be maximized or minimized simultaneously for all vertex classes of a partition. For instance, for a graph G, what is the minimum of $\max\{e(V_1), e(V_2)\}$ over all partitions $V(G) = V_1 \cup V_2$?
After discussing some extremal results for Max Cut, and related algorithms, we present some results on judicious partitions for graphs and hypergraphs and some open problems.

Logarithmic Sobolev Constants & Average Conductance of Balanced Matroids

Jung-Bae Son
Dept. of Computer Science
Edinburgh University

The notion of balanced matroids was first coined by Feder and Mihail in 1992. They used conductance, as defined by Jerrum and Sinclair, and canonical paths techniques to show that the random walk on the bases exchange graph of balanced matroids is rapidly mixing.
We use two recent techniques, Kannan and Lovász's average conductance and new lower bounds by Houdré on logarithmic Sobolev constants, to improve Feder and Mihail's bounds for certain balanced matroids, namely regular matroids with a constant number of parallel elements.
Joint work with Ravi Montenegro (Yale University).

Optimal myopic algorithms for random 3-SAT

Gregory Sorkin
Mathematical Sciences Dept.
IBM T.J. Watson Research Center

3-SAT is a canonical NP-complete problem: satisfiable and unsatisifiable instances cannot generally be distinguished in polynomial time. However, random 3-SAT formulas show a phase transition: sparse instances are almost always satisfiable, and dense ones almost always unsatisfiable.
Proofs of the satisfiability of sparse instances have come from analyzing simple heuristics: the better the heuristic analyzed, the denser the instances that can be proved satisfiable with high probability. To date, the useful heuristics have all been simple extensions of unit-clause propagation, all expressible within a common framework, and analyzable in a uniform manner by employing differential equations.
Here, we determine optimal algorithms expressible in that framework, establishing an improved density bound. We extend the analysis via differential equations, and make extensive use of a new optimization problem we call "max-density multiple-choice knapsack". The structure of optimal knapsack solutions elegantly characterizes the choices made by an optimal algorithm.
Joint work with Dimitris Achlioptas.

A new performance measure for stochastic scheduling

Angelika Steger
Institut für Informatik
Technische Universität München

A common approach in stochastic scheduling is to minimize the expectation of the objective function under consideration (e.g. makespan or sum of completion times). Unfortunately, the expectation does not take into account the variance of the distributions. It is therefore easy to come up with examples for which there are clearly better strategies than those which minimize the expectation. We therefore propose to use a different performance measure, namely, the expectation of the competitive ratio. We also show that in the case of exponentially distributed random variables the strategy "shortest expected processing times first" has a constant performance ratio with respect to the sum of the completion time.
Joint work with Mark Scharbrodt and Thomas Schickinger.

Can Entropy Characterize Performance of Online Algorithms?

Eli Upfal
Computer Science Department
Brown University

Viewing online problems with stochastic input as iterative gambling games, we explore the relation between the entropy of the input sequence and the performance of the best online algorithm for that problem. We present both positive and negative results, showing that entropy is a good performance characterizer for list accessing and prefetching problems, but a poor characterizer for online caching. The motivation for this work are advanced system and architecture designs which allow the operating system to dynamically allocate resources to online protocols such as prefetching and caching. To utilize these features the operating system needs to identify data streams that can benefit from more resources. This question is not addressed by the standard online competitive analysis.

Perfect Simulation for Quenched Disordered Systems

David B. Wilson
Microsoft Research
Redmond

This is a two-part talk; in the first part we explain the read-once CFTP method of perfect simulation, and in the second part we report on an application of perfect simulation to the study of quenched disordered systems, which is joint work with Gilles Schaeffer. In computer science, statistics, and physics it is often desirable to generate random configurations drawn from some probability distribution. One prevalent method for doing this is to construct a Markov chain whose stationary distribution is the desired distribution, and then run the Markov chain for “a long time”. There are a variety of methods for determining how long to run the Markov chain. Coupling from the past (CFTP) (due to Jim Propp and the speaker) is a method whereby the computer determines on its own how long to run the Markov chain, and returns a sample drawn exactly according to the stationary distribution of the Markov chain. Those acquainted with the CFTP method of perfect simulation will recall that the algorithm sometimes needs to re-use old random coins, and that flipping fresh random coins at these times will introduce bias. Read-once CFTP is a variation of CFTP that only reads random coins once. The second part of the talk, where we discuss the use of perfect simulation to study statistical mechanical systems with quenched disorder, is also partly expository, since we explain the use of additional techniques that members of the audience may find useful in other contexts.
Joint work with Gilles Schaeffer.