05201 Abstracts Collection
Design and Analysis of Randomized and Approximation Algorithms
— Dagstuhl Seminar —

Martin Dyer[1], Mark Jerrum[2] and Marek Karpinski[3]

[1] Univ. of Leeds, GB
dyer@comp.leeds.ac.uk
[2] Univ. of Edinburgh, GB
mrj@inf.ed.ac.uk
[3] Univ. Bonn, DE
marek@cs.uni-bonn.de

Abstract. From 15.05.05 to 20.05.05, the Dagstuhl Seminar 05201 "Design and Analysis of Randomized and Approximation Algorithms" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Randomized Algorithms, Approximation Algorithms, Optimization Problems, Measurement Problems, Decentralized Networks

05201 Overview – Design and Analysis of Randomized and Approximation Algorithms

Most computational tasks today that arise in realistic scenarios are intractable, at least if one insists on exact solutions delivered within a strict deadline. Two important means for surmounting that intractability barier are randomized and approximate computations. It is an intersting artifact that these two notions of computation go hand-in-hand.

The Seminar was concerned with the newest developments in the design and analysis of randomized and approximation algorithms. The main focus of the workshop was on three specific topics: approximation algorithms for optimization problems, approximation algorithms for measurement problems, and decentralized networks as well as various interactions between them. Here, some new broadly applicable techniques have emerged recently for designing efficient approximation algorithms for various optimization and measurement problems.

Dagstuhl Seminar Proceedings 05201
Design and Analysis of Randomized and Approximation Algorithms
http://drops.dagstuhl.de/opus/volltexte/2005/319

This workshop has addressed the above topics and also some new fundamental insights into the design techniques.

The 35 lectures delivered at this workshop covered a wide body of research in the above areas. The Program of the meeting and Abstracts of all talks are listed in the subsequent sections of this report.

The meeting was hold in a very informal and stimulating atmosphere. Thanks to everyone who made it a very interesting and enjoyable event.

Martin Dyer
Mark Jerrum
Marek Karpinski

Acknowlegement. We thank Annette Beyer, and Angelika Mueller for their continuous support and help in organizing this workshop.

Joint work of: Dyer, Martin; Jerrum, Mark; Karpinski, Marek

Informal Session on Algorithmic Game Theory

In this informal session which was held on thursday May 19th, we gave a brief survey on our current work on Market Equilibria.

In Fisher markets we are given a set of buyers B and a set of goods A. Each buyer $b \in B$ has an initial endowment of money e_b, the amount of each good can be normalized to 1. Buyers $b \in B$ have utility functions $u_b\colon \mathbf{R}_+^A \to \mathbf{R}$. At given prices $p_j, j \in A$ each buyer b wants to optimize (i.e. maximize) her total utility:

$$\max\ u_b(x) \text{ s.t. } x \in [0,1]^A,\ \sum_{j \in A} p_j \cdot x_j \leq e_b$$

A price vector $p = (p_j | j \in A)$ is called Market Equilibrium iff there exists an assignment $x = (x_{bj})_{b \in B, j \in A}$ of goods to buyers such that (1) the market clears: $\forall j \in A\ \sum_b x_{bj} = 1$ and (2) $\forall b \in B\ u_b(x_b) = u_b(x_{bj} | j \in A)$ is an optimum solution to buyer b's utility maximization problem at given prices p. The price vector p is calles ϵ-approximate market equilibrium if

$$\sum_{j \in A} |p_j - d_j| \quad \leq \quad \epsilon \cdot \sum_{j \in A} p_j$$

where d_j is the total demand for good j at price p (i.e. $d_j = \sum_{b \in B} x^\star_{bj}$ where $x_b \in [0,1]^A$ is an opt. solution to buyer b's optimization problem at price p).

(a) For utilities $u_b(x) = \sum_{j \in A}(a_{bj} \log(x_j) + c_{bj})$ we give a PTAS based on a simple greedy approach. This carries over to the case of non-additive separable utilities $u_b(x)$ with the property that the partial derivatives satisfy the equation

$$\frac{\partial}{\partial x_j} u_b(x) = \frac{1}{\alpha} \cdot \frac{\partial}{\partial x_j} u_b\left(\frac{x}{\alpha}\right)$$

for every $x > 0$ and $\alpha \geq 1$.

(b) We describe a general framework for normal demand markets to obtain a PTAS for the Arrow-Debreu case (exchange markets - buyers do not have money but initial amounts of goods) whenever a polynomial time algorithm for the Fisher case is given.

Joint work of: Hauptmann, Mathias; Langguth, Johannes

Improved Approximation Algorithms for Metric Maximum ATSP and Maximum 3-Cycle Cover Problems

Markus Bläser (ETH Zürich, CH)

We consider an APX-hard variant (DATSP) and an APX-hard relaxation (DCC) of the classical traveling salesman problem. DATSP is the following problem: Given an edge-weighted complete loopless directed graph G such that the edge weights fulfill the triangle inequality, find a maximum weight Hamiltonian tour of G. We present a 31/40-approximation algorithm for DATSP with polynomial running time.

DCC is the following problem: Given an edge-weighted complete loopless directed graph, compute a collection of node-disjoint cycles, each of length at least three, whose weight is maximum among all such collections. We present a 3/4-approximation algorithm for this problem with polynomial running time. In both cases, we improve on the previous best approximation performances. The results are obtained via a new decomposition technique for the fractional solution of an LP formulation of DCC.

Joint work of: Bläser, Markus; Shankar Ram, L.; Sviridenko, Maxim

Hypergraph Independent Sets and Colorings

Magnus Bordewich (University of Leeds, GB)

We consider the mixing time of Glauber dynamics for sampling independent sets in hypergraphs. A simple analysis based on path coupling as in the graph case gives weak results.

Intuitively large edges should mean that we move away from bad-case situations, and hence you expect to couple before diverging, but only at a stopping time not in one step. We prove a result that path coupling can be made to work in this case (similar results by Hayes and Vigoda are in FOCS 04) and apply this.

We obtain results for sampling colourings in hypergraphs also.

Joint work of: Bordewich, Magnus; Dyer, Martin; Karpinski, Marek

Proof of the local REM conjecture for number partitioning

Christian Borgs (Microsoft Research - Seattle, USA)

The number partitioning problem is a classical combinatorial optimization problem: Given n numbers or weights, one is faced with the problem of partitioning this set of numbers into two subsets to mininize the discrepancy, defined as the absolute value of the difference in the total weights of the two subsets.

Here we consider random instances of this problem where the n numbers are i.i.d. random variables, and we study the distribution of the discrepancies and the correlations between partitions with similar discrepancy. In spite of the fact that the discrepancies of the 2^{n-1} possible partitions are clearly correlated, a surprising recent conjecture states that the discrepancies near any given threshold become asymptotically independent, and that the partitions corresponding to these discrepancies become uncorrelated. In other words, the conjecture claims that near any fixed threshold, the cost function of the number partitioning problem behaves asymptotically like a random cost function.

In this talk, I describe our recent proof of this conjecture.

Joint work of: Borgs, Christian; Chayes, Jennifer; Mertens, Stephan; Nair, Chandra

Aggregating Inconsistent Information: Ranking and Clustering

Moses Charikar (Princeton University, USA)

A ranking of n web pages is to be chosen from outputs of k search engines. How do we choose one ranking minimizing the "disagreement" with the k rankings ? A clustering of n objects is to be chosen from outputs of k clustering algorithms. How do we choose one clustering minimizing the "disagreement" with the k clusterings?

These information aggregation problems date back to 1785, when the French philosopher Condorcet considered voting systems where each voter chooses a full ranking of a set of candidates. Recently, applications in computer science have motivated the study of their algorithmic and complexity aspects. I will show that for both these problems, we can obtain improved algorithms using essentially the same, remarkably simple principle.

Furthermore, this also applies to and yields improvements for related graph theoretic optimization problems, known as the minimum feedback arc set in tournaments and the correlation clustering in complete graphs.

Additionally, I will show that the problem of finding a minimum feedback arc set in tournaments has no poly-time algorithm, assuming NP is not contained in BPP, almost settling a long-standing conjecture of Bang- Jensen and Thomassen.

This is joint work with Nir Ailon and Alantha Newman.

Joint work of: Ailon, Nir; Charikar, Moses; Newman, Alantha

On Minimum Concave Cost Multicommodity Flow (a.k.a. Buy-at-Bulk network design)

Moses Charikar (Princeton University, USA)

We study the multicommodity buy-at-bulk network design problem where the goal is to buy capacity on edges of a network so as to enable the demands between a given set of source-sink pairs to be routed - the objective is to minimize the cost of such a solution. The key aspect of this problem is that the cost of an edge in the graph is a concave monotone function of the flow cross the edge and hence exhibits economy of scale - it pays to aggregate flow paths as much as possible. In the non-uniform case, each edge has its own cost function, possibly different from other edges.

Special cases of this problem have been studied extensively. We present the first non- trivial approximation algorithm for the general case. Our algorithm is an extremely simple randomized greedy algorithm, involving not much more than shortest path computations.

We achieve an approximation guarantee of exp (O(sqrt(log(n)log(log(n)))) where n is the total demand in the graph.

This is joint work with Adriana Karagiozova

Joint work of: Charikar, Moses; Karagiozova, Adriana

Sorting with Restricted Set of Allowed Comparisons

Moses Charikar (Princeton University, USA)

We consider the following variant of Sorting: Given n Items and a graph $G = (V, E)$ on $V = \{1, \ldots, n\}$, we are allowed to compare two items if and only if the corresponding edge exists in graph G.

Question: How many comparisons do we need ?

(1) Are there graphs for which more than $n \log(n)$ (ie. $\omega(n \cdot \log(n))$) are required ?

(2) Is there an algorithm that takes $n^{2-\epsilon}$ comparisons for **any** graphs ?

The Spread of Viruses on Scale-Free Graphs

Jennifer Chayes (Microsoft Research - Seattle, USA)

We analyze the contact process on random graphs generated according to a preferential attachment scheme as a model for the spread of viruses on scale-free graphs. We show that any virus with a positive rate of spread from a node to its neighbors has a non-vanishing chance of becoming epidemic. As a function of the rate of spread from a node to its neighbors, we determine both the epidemic probability starting from a typical vertex, and the average epidemic probability – which is dominated by atypical vertices. In the process of this work, we develop a new representation for the mathematical analysis of scale-free graphs.

Joint work of: Berger, Noam; Borgs, Christian; Chayes, Jennifer; Saberi, Amin

A Spectral Heuristic for Bisecting Random Graphs

Amin Coja-Oghlan (HU Berlin, D)

The minimum bisection problem is to partition the vertices of a graph into two classes of equal size so as to minimize the number of crossing edges. The problem is NP-hard in the worst case. We analyze a spectral heuristic for this problem on random graphs with a planted bisection, extending the work of Boppana (FOCS 1987).

Keywords: Minimum bisection, random graphs, spectral techniques

Sampling Integer Network Flows

Mary Cryan (University of Edinburgh, GB)

We consider the problem of approximately counting integral flows in a network. We show that there exists an FPRAS based on volume estimation if all capacities in the network are sufficiently large, generalizing a result of Dyer, Kannan and Mount (1997). Our proof relies on the properties of the Maximum Spanning Tree in the network (using the capacities as weights) to show that the full-dimensional Flow Polytope is well-rounded.

The fpras for integral flows can be applied to the problem of cell-bounded contingency tables, whenever the parameters of this problem are sufficiently large.

We can also use a combination of dynamic programming and volume estimation to approximately count cell-bounded contingency tables when the number of rows is constant, regardless of the sizes of the cell bounds.

Joint work of: Cryan, Mary; Dyer, Martin; Randall, Dana

Estimating the weight of metric minimum spanning trees in sublinear-time

Artur Czumaj (NJIT - Newark, USA)

In this paper we present a sublinear time $(1 + \epsilon)$-approximation randomized algorithm to estimate the weight of the minimum spanning tree of an n-point metric space. The running time of the algorithm is $\widetilde{O}(n/\epsilon^{O(1)})$. Since the full description of an n-point metric space is of size $\Theta(n^2)$, the complexity of our algorithm is *sublinear* with respect to the input size. Our algorithm is almost optimal as it is not possible to approximate in $o(n)$ time the weight of the minimum spanning tree to within any factor. Furthermore, it has been previously shown that no $o(n^2)$ algorithm exists that returns a spanning tree whose weight is within a constant times the optimum.

Keywords: Minimum spanning tree, sublinear-time algorithms, approximation algorithms

Joint work of: Czumaj, Artur; Sohler, Christian

See also: Appeared in Proceedings of the 36th ACM Symposium on Theory of Computing (STOC'04), pages 175 - 183, Chicago, IL, June 13 - 15, 2004. ACM Press, New York, NY, 2004

Sampling Regular Graphs and a Peer-to-Peer Network

Martin E. Dyer (University of Leeds, GB)

The problem of uniformly sampling graphs on a fixed number of vertices with a given degree sequence has been well studied, particularly in the case of sampling d-regular graphs on n vertices, where $d = d(n)$ may grow with n. We consider generation by a simple Markov chain on d-regular graphs with n vertices, and show that this has mixing time to the uniform distribution bounded above by a polynomial in n and d.

A related Markov chain on d-regular graphs with a varying number of vertices is introduced, for even constants d. This is used to model a certain peer-to-peer network structure proposed in the literature, where we assume that clients arrive and depart according to a size-dependent Poisson process. We prove that this process has mixing time to its equilibrium distribution which is bounded above by a polynomial in N, the expected number of clients in the system, provided certain assumptions are met about the rate of arrival and departure of clients.

On the average case performance of some greedy approximation algorithms for the uncapacitated facility location problem

Alan M. Frieze (CMU - Pittsburgh, USA)

In combinatorial optimization, a popular approach to NP-hard problems is the design of approximation algorithms. These algorithms typically run in polynomial time and are guaranteed to produce a solution which is within a known multiplicative factor of optimal.

Unfortunately, the known factor is often known to be large in pathological instances. Conventional wisdom holds that in practice, approximation algorithms will produce solutions closer to optimal than their proven guarantees. In this paper we use the rigorous-analysis-of-heuristics framework to investigate this conventional wisdom.

We analyze the performance of 3 related approximation algorithms for the uncapacitated facility location problem, when each is applied to instances created by placing n points uniformly in the unit square. We find that whp these 3 algorithms do not find asymptotically optimal solutions and that also whp a simple plane partitioning algorithm does.

Joint work of: Frieze, Alan M.; Flaxman, Abie; Vera, Juan

On Packing Squares with Resource Augmentation: Maximizing the Profit

Olga Gerber (Universität Kiel, D)

We consider the problem of packing squares with profits into a bounded square region so as to maximize their total profit. More specifically, given a set L of n squares with positive profits, it is required to pack a subset of them into a unit size square region $[0,1] \times [0,1]$ so that the total profit of the squares packed is maximized. For any given positive accuracy $\epsilon > 0$, we present an algorithm that outputs a packing of a subset of L in the augmented square region $[1+\epsilon \times [1+\epsilon]$ with profit value at least $(1-\epsilon)OPT(L)$, where $OPT(L)$ is the maximum profit that can be achieved by packing a subset of L in the unit square.

The running time of the algorithm is polynomial in n for fixed ϵ.

Joint work of: Gerber, Olga; Fishkin, Aleksei; Jansen, Klaus; Solis-Oba, Roberto

Algorithmic Aspects of the Regularity Lemma for Sparse Graphs

Stefanie Gerke (ETH Zürich, CH)

Szemeredi's regularity lemma has many applications, for example in extremal graph theory, in property testing, and in the development of approximation algorithms. By the nature of Szemeredi's lemma, most applications have been to dense instances. In 1997 Kohayakawa and Rödl introduced a version of the regularity lemma that can be applied to a class of sparse graphs, namely so-called bounded graphs. We show how to find a partition guaranteed by this version of the regularity lemma in time $n^{O(1/d)}$ where d is the density of the graph. This implies for example that for bounded graphs, given an $\epsilon > 0$ one may approximate Max Cut up to an additive error of ϵdn^2 in time $n^{O(1/d)}$.

Joint work of: Gerke, Stefanie; Huber, Manuel; Steger, Angelika

Sampling Anti-ferromagnetic Potts configurations on the integer lattice

Leslie Ann Goldberg (University of Warwick, GB)

We study Glauber dynamics for the anti-ferromagnetic Potts model on Z^2. There are two parameters — the number of spins, and the temperature, which determines the weight of a monochromatic edge.

The goal is to determine for which parameters the dynamics is rapidly mixing.

I'll describe what is known about this problem, including some work with Markus Jalsenius, Russ Martin, and Mike Paterson.

The paper http://www.dcs.warwick.ac.uk/people/academic/Leslie.Goldberg/papers/spatial.ps by Goldberg, Paterson, and Martin contains results for the zero temperature case. Work in progress with Jalsenius extends this to other temperatures.

Keywords: Rapid mixing, colourings

Full Paper:
www.dcs.warwick.ac.uk/people/academic/Leslie.Goldberg/papers/spatial.ps

Simple Sampling Solutions for Stochastic Optimization

Anupam Gupta (CMU - Pittsburgh, USA)

Stochastic optimization involves making decisions under uncertain conditions.

We consider problems in the model of "two stage stochastic optimization with recourse", where this can be easily paraphrized as follows:

(1) On monday, we are given a distribution π from which the clients will be drawn on tuesday. We are allowed to buy some elements T_0 in anticipation of this demand, where the cost of element e is c_e.

(2) On tuesday, a set S of clients appears, which is drawn from the distribution π.

One is allowed to buy more elements T_1 to form a feasible solution $(T_0 \cup T_1)$ to the clients in S. However, the elements now cost $\sigma \cdot c_e$ (for $\sigma > 1$).

We present simple solutions for problems such as Steiner Tree, Vertex Cover, Facility Location in this model. In fact, we give a general way to convert algorithms for the deterministic versions of these problems to stochastic versions of the problems.

Joint work of: Gupta, Anupam; Pal, Martin; Ravi, R.; Sinha, Amitabh

Bandwidth Allocation

Anupam Gupta (CMU - Pittsburgh, USA)

We are given a graph $G = (V, E)$, edge costs $c\colon E \to \mathbb{R}$ and a subset $S \subseteq V$ such that each $s \in S$ wants to communicate at rate 1. We ask for a valid traffic matrix (d_{ij}) where d_{ij} denotes the traffic demand between node i and node j. The conditions $\forall i \ \sum_j d_{ij} \leq 1$ define the fractional matching polytope P.

We want to allocate bandwidth x_e on edges e, hence the resulting optimization problem is the following:

Minimize $\sum_e c_e x_e$ such that for all $\bar{d} \in P$ there exists a flow f_d such that (a) f_d satisfies $\bar{d}$ and (b) f_d respects the capacities x_e.

Question: Can we do it in polynomial time ?

It is known that there is a 2-approximation algorithm for this problem. Furthermore, if the graph is directed then the problem becomes coNP-hard (Gupta, Kleinberg, Kumar, et al., STOC'01).

Every 2-CSP Allows Nontrivial Approximation

Johan Hastad (KTH Stockholm, S)

We use semidefinite programming to prove that any constraint satisfaction problem in two variables over any domain allows an efficient approximation algorithm that does provably better than picking a random assignment. To be more precise assume that each variable can take values in $[d]$ and that each constraint rejects t out of the d^2 possible input pairs. Then, for some universal constant c, we can, in probabilistic polynomial time, find an assignment whose objective value is, on expectation, within a factor $(1 - \frac{t}{d^2}(1 - \frac{c}{d^2 \log d}))$ of optimal.

A general lower bound for mixing of single-site dynamics on graphs

Thomas Hayes (Univ. California - Berkeley, USA)

We prove that any Markov chain that performs local, reversible updates on randomly chosen vertices of a bounded-degree graph necessarily has mixing time at least Omega(n log n), where n is the number of vertices. Our bound applies to the so-called "Glauber dynamics" that has been used extensively in algorithms for the Ising model, independent sets, graph colorings and other structures in computer science and statistical physics, and demonstrates that many of these algorithms are optimal up to constant factors within their class. Previously, no super-linear lower bound was known for this class of algorithms. Though widely conjectured, such a bound had been proved previously only in very restricted circumstances, such as for the empty graph and the path.

The main ingredients in the proof are "disagreement percolation," which bounds the rate at which information can flow between vertices as a function of their distance; "complete monotonicity," which implies that, for suitable initial conditions, the probability of a vertex having its initial value decreases monotonically with time; and coupling. We also show that the assumption of bounded degree is necessary by giving a family of dynamics on graphs of unbounded degree with mixing time o(n log n).

Joint work with Alistair Sinclair.

Keywords: Mixing time, lower bound, Glauber dynamics, graphs, Ising model, graph colorings, independent sets

Joint work of: Hayes, Thomas; Sinclair, Alistair

Building Scalable and Robust Peer-to-Peer Overlay Networks for Broadcasting using Network Coding

Kamal Jain (Microsoft Research - Seattle, USA)

We propose a scheme for building peer-to-peer overlay networks for broadcasting using network coding. The scheme addresses many practical issues such as scalability, robustness, constraints on bandwidth, and locality of decisions. We analyze the system theoretically and prove near optimal bounds on the parameters defining robustness and scalability.

As a result we show that the effects of failures are contained locally, allowing the network to grow exponentially with server load. We also argue that adversarial failures are no more harmful than random failures.

Keywords: Network coding; random graphs; peer-to-peer

Joint work of: Jain, Kamal; Lovasz, Laszlo ; Chou, Philip A.

Approximation algorithms for scheduling malleable tasks with precedence constraints

Klaus Jansen (Universität Kiel, D)

In this talk we present approximation algorithms for scheduling malleable tasks with precedence constraints. The best previous approximation algorithm (that works in two phases) by Lepére, Trystram and Woeginger has a ratio $3+\sqrt{5} \approx 5.236$. In the first phase a discrete time-cost tradeoff problem is solved approximately, and in the second phase a variant of the list scheduling algorithm is used. In phase one a rounding parameter $\rho = 1/2$ and in phase two an allotment parameter $\mu = (3m - \sqrt{5m2 - 4m})/2$ is employed, respectively. We study the influence of the rounding parameter ρ to the second phase. With setting $\rho \neq 1/2$ we are able to obtain a better approximation algorithm with a ratio of $(3+\sqrt{5})/2 + \sqrt{2(\sqrt{5}+1)} \approx 5.162$. Furthermore, we study the linear program relaxation and rounding technique in the first phase more carefully and its influence to the second phase. As a result we develop an improved approximation algorithm with a ratio of $100/43 + 100(\sqrt{4349} - 7)/2451 \approx 4.730598$. In addition we study another model of the malleable tasks (where the speed-up function is concave). Instead of solving the discrete time-cost tradeoff problem, we solve here a piecewise linear program and use a new rounding technique. In this model we obtain an approximation algorithm with a ratio of $100/63 + 100(\sqrt{6469} + 13)/5481 \approx 3.291919$.

Keywords: Approximation algorithms, scheduling malleable tasks

Joint work of: Jansen, Klaus; Zhang, Hu

The complexity of the ferromagnetic Ising model with local fields

Mark Jerrum (University of Edinburgh, GB)

After reviewing the arcane world of complexity classes associated with approximate counting problems in general, we turn to consider a specific problem: namely, approximating the partition function of the ferromagnetic Ising model with varying interaction energies and local external magnetic fields. Jerrum and Sinclair provided a fully polynomial randomised approximation scheme for the case in which the system is consistent, in the sense that the local external fields all favour the same spin. We characterise the complexity of the general problem by showing that it is equivalent in complexity to the problem of approximately counting independent sets in a bipartite graph; thus it is complete in a logically-defined subclass of $\sharp P$ previously studied by Dyer, Goldberg, Greenhill and Jerrum. In contrast, we show that approximating the partition function of the q-state Potts model with q>2 is as hard as approximately solving any counting problem in $\sharp P$.

Keywords: Approximate counting, computational complexity, hard-core gas model, independent sets, Ising model

Joint work of: Goldberg, Leslie; Jerrum, Mark

Approximating Metric Bisection and Related Partitioning Problems

Marek Karpinski (Universität Bonn, D)

We design polynomial time approximation schemes (PTASs) for the Metric Bisection problem, i.e. the problem of dividing a given finite metric space into two halves so as to minimize the sum of distances across the cut. This settles the status of that problem which was open for some time now. Our method generalizes for metric Partitioning problems with arbitrary size constraints including Metric MIN-MULTIWAY-k-CUT problems.

The solution depends on a development of certain hybrid placement methods combined with a new method of linearizing smooth metric polynomials programs. We discuss also several possible directions for improving the efficiency and the sample complexity of our PTASs and the connection to the newly discovered efficient core metric and quasimetric decomposition methods.

Keywords: Bisection

Joint work of: Karpinski, Marek; de la Vega, W.F.; Kenyon, Claire

Strong Inapproximability Results for Metric TSP

Marek Karpinski (Universität Bonn, D)

We review some explicit inapproximability results of Papadimitriou-Vempala, 2002, and Engebretsen-Karpinski, 2004, for metric TSP and bounded metric TSP problems, respectively.

We notice that the inapproximability results for metric TSP with integer distances between one and eight almost match the general metric TSP bound, and that all known up to date methods are 'local' in nature. We ask for existence of 'non-local' methods for proving stronger inapproximability results for metric TSP.

Perhaps specially tailored direct PCP techniques could improve known inapproximability results by orders of magnitude.

On the approximability of non-Boolean Max CSP

Andrei Krokhin (University of Durham, GB)

In a maximum constraint satisfaction problem (Max CSP), one is given a finite collection of (possibly weighted) constraints on overlapping sets of variables, and the goal is to assign values from a given finite domain to the variables so as to maximize the number (or the total weight) of satisfied constraints. This problem is NP-hard in general, and, therefore, it is natural to study how restricting the allowed types of constraints affects the approximability of the problem.

It is known that every Boolean Max CSP problem with a finite set of allowed constraint types is either solvable exactly in polynomial time or else APX-complete (and hence can have no PTAS unless P=NP). It has been an open problem for several years whether this result can be extended to non-Boolean Max CSP. In this talk, we will discuss recent progress made on this classification problem.

Keywords: Constraint satisfaction problem complexity approximability

Joint work of: Krokhin, Andrei; Jonsson, Peter; Klasson, Mikael

The Encoding Complexity of Network Coding

Michael Langberg (CalTech - Pasadena, USA)

How many edges need to perform encoding in the Multicast Network Coding problem ? It turns out that this problem is NP-hard to approximate within ration $|E|^{1-\epsilon}$ (multiplicative/additive).

We study universal bounds on the number of encoding edges needed to transmit h messages to k destinations.

In acyclic networks this number is at most $O(h^3 \cdot k^2)$ and (globally) cannot be less that $\Omega(h^2 k)$. In networks with cycles we bound this number between $O((FES+1)h^3k^2)$ and $\Omega(FES \cdot h)$. Here FES is the Minimum feedback edge set of the graph.

Joint work of: Langberg, Michael; Sprintson, Alex; Bruck, Jehoshua

Approximating the maximum clique minor and some subgraph homeomorphism problems

Andrzej Lingas (Lund University, S)

We consider the "minor" and "homeomorphic" analogues of the maximum clique problem, i.e., the problems of determining the largest h such that the input graph (on n vertices) has a minor isomorphic to K_h or a subgraph homeomorphic to K_h, respectively, as well as the problem of finding the corresponding subgraphs. We term them as the *maximum clique minor* problem and the *maximum homeomorphic clique* problem, respectively. We observe that a known result of Kostochka and Thomason supplies an $O(\sqrt{n})$ bound on the approximation factor for the maximum clique minor problem achievable in polynomial time. We also provide an independent proof of nearly the same approximation factor with explicit polynomial-time estimation, by exploiting the minor separator theorem of Plotkin *et al.* Next, we show that another known result of Bollobás and Thomason and of Komlós and Szemerédi provides an $O(\sqrt{n})$ bound on the approximation factor for the maximum homeomorphic clique achievable in polynomial time. On the other hand, we show an $\Omega(n^{1/2-O(1/(\log n)^{\gamma})})$ lower bound (for some constant γ, unless $\mathcal{NP} \subseteq \mathrm{ZPTIME}(2^{(\log n)^{O(1)}})$) on the best approximation factor achievable efficiently for the maximum homeomorphic clique problem, nearly matching our upper bound.

Finally, we derive an interesting trade off between approximability and subexponential time for the problem of subgraph homeomorphism where the guest graph has maximum degree not exceeding three and low treewidth. In particular, we observe that by a result of Fellows and Langston for any graph G on n vertices and a positive integer q not exceeding n, one can produce either an n/q approximation to the longest cycle problem in polynomial time, or find an optimal longest cycle of G in time $2^{O(q \log q + \log n)}$.

Joint work of: Lingas, Andrzej; Alon, Noga; Wahlen, Martin

Distributed Selfish Load Balancing

Russell Martin (University of Warwick, GB)

Suppose a set of m tasks are to be shared as equally as possible amongst a set of n resources. A game-theoretic mechanism to find a suitable allocation is to associate each task with a "selfish agent", and require each agent to select a resource, with the cost of a resource being the number of agents to select it. Agents would then be expected to migrate from overloaded resources until the allocation becomes balanced.

Recent work has studied the question of how this can take place within a distributed setting in which agents migrate selfishly without any centralized control. We discuss a natural protocol for the agents which combines the following desireable features:

It can be implemented in a strongly distributed setting, using no central control, and has good convergence properties. We show using a martingale technique that the process converges in expected time $O(\log\log m + n^4)$.

We also give a lower bound of $\Theta(\max\{\log\log m, n\})$ for the convergence time, as well as an exponential lower bound (in n) for a variation of this protocol that allows the agents to migrate even if they do not strictly improve their situation.

Joint work of: Martin, Russell; Berenbrink, Petra; Friedetzky, Tom; Goldberg, Leslie Ann; Goldberg, Paul; Hu, Zengjian

Succinct Constructions of k-Wise (Almost) Independent Permutations

Moni Naor (Weizmann Inst. - Rehovot, IL)

Constructions of k-wise almost independent permutations have been receiving a growing amount of attention in recent years. However, unlike the case of k-wise independent functions, the size of previously constructed families of such permutations is far from optimal.

In this talk we will describe a method for reducing the size of families given by previous constructions. Our method relies on pseudorandom generators for space-bounded computations. More specifically, we show that a weaker derandomization tool suffices: all we need is a generator that produces "pseudorandom walks" on undirected graphs with a consistent labelling.

Such generators with sufficiently good parameters are implied by the proof that undirected connectivity is in logspace of Reingold, and made explicit by Reingold, Trevisan and Vadhan.

With this approach we are able to obtain families of k-wise almost independent permutations of size that is optimal up to polynomial factor. More precisely, if the distance from uniform for any k tuple should be at most delta, then the size of the description of a permutation in the family is O(kn + log 1/delta).

Keywords: Derandomization, permutations, shuffling, random walks

Joint work of: Kaplan, Eyal; Naor, Moni; Reingold, Omer

LP based Solutions for Stochastic Optimization

Ramamoorthi Ravi (Carnegie Mellon University - Pittsburgh, USA)

Real-world networks often need to be designed under uncertainty, with only partial information and predictions of demand available at the outset of the design process. The field of stochastic optimization deals with such problems where the forecasts are specified in terms of probability distributions of future data.

In this talk, we apply linear programming based methods for approximating stochastic optimization problems on networks. For example, we look at stochastic models where the cost of the elements is correlated to the set of realized demands, and risk-averse models where upper bounds are placed on the amount spent in each of the stages. These generalized models require new techniques, and our solutions are based on a combination of the primal-dual method truncated based on optimal LP relaxation values, followed by a tree-rounding stage. We use these to give constant-factor approximation algorithms for the stochastic Steiner tree and single sink network design problems in these generalized models.

This talk is based on a paper co-authored with Anupam Gupta, CMU, and Amitabh Sinha, University of Michigan, that appeared in FOCS 2004.

Approximating two-stage stochastic problems II: LP-based methods

Ramamoorthi Ravi (Carnegie Mellon University - Pittsburgh, USA)

Real-world networks often need to be designed under uncertainty, with only partial information and predictions of demand available at the outset of the design process. The field of stochastic optimization deals with such problems where the forecasts are specified in terms of probability distributions of future data.

In this talk, we apply linear programming based methods for approximating stochastic optimization problems on networks. For example, we look at stochastic models where the cost of the elements is correlated to the set of realized demands, and risk-averse models where upper bounds are placed on the amount spent in each of the stages. These generalized models require new techniques, and our solutions are based on a combination of the primal-dual method truncated based on optimal LP relaxation values, followed by a tree-rounding stage. We use these to give constant-factor approximation algorithms for the stochastic Steiner tree and single sink network design problems in these generalized models.

This talk is based on a paper co-authored with Anupam Gupta, CMU, and Amitabh Sinha, University of Michigan, that appeared in FOCS 2004.

Computing (Correlated) Equilibria in Multi-Player Games

Tim Roughgarden (Stanford University, USA)

This talk will survey recent research (joint with Christos Papadimitriou) on the complexity of computing equilibria - particularly correlated equilibria - in games with a large number of players. Such games, in order to be computationally meaningful, must be presented in some succinct, game-specific way. I will discuss a general algorithmic framework for identifying the complexity of optimizing over the correlated equilibria of a compactly represented multi-player game, and applications of this paradigm to symmetric games, graphical games, and congestion games. Consequences include (on the positive side) the first polynomial-time algorithm for this problem in symmetric games and (on the negative side) the first complexity-theoretic justification of the standard restrictions on the topology of a graphical game.

Efficiency of Equilibria in Network Games

Tim Roughgarden (Stanford University, USA)

Consider network games where $G = (V, E)$ is a graph with edge costs c_e, $e \in E$ and players $i \in N = \{1, \ldots, k\}$ such that each player i corresponds to a pair (s_i, t_i) of vertices from G. The strategies of player i are the $s_i - t_i$-paths in G. A strategy profile in this game is a network $\bigcup_{i=1}^{k} P_i$ where P_i is the strategy played by player i $(1 \leq i \leq k)$. Players using edge e split e equally, hence the cost to player i is

$$\sum_{e \in P_i} \frac{c_e}{\sharp \text{ players using } e}.$$

The **Price of Stability** is defined as the ratio $\frac{\text{Best Nash Eq.}}{OPT}$.

It is known that the Price of Stability is $\mathcal{H}_k$ in directed graphs (Anshelvich et al. FOCS'04).

Open: What is the Price of Stability for undirected graphs ? The best known lower bound is 2.

The Query Complexity of Local Search

Miklos Santha (Université Paris Sud, F)

We prove that the deterministic, randomized and quantum query complexity of local search are polynomially related in every graph. In this problem, given a graph G, and an oracle function $f\colon V \to \mathbb{Z}$, we look for a local minimum, i.e. a vertex v, such that for every w $(\{v, w\} \in E)$ $f(v) \leq f(w)$.

Joint work of: Santha, Miklos; Szegedy, Mario

Asynchronous randomized automata: how does randomness affect decentralized algorithms execution

Nicolas Schabanel (ENS - Lyon, F)

In this paper we consider a certain type of decentralized and distributed algorithms and study how much asynchronism affect their execution. We propose a probabilistic analysis of the fully asynchronous behavior (i.e., two cells are never simultaneously updated, as in a continuous time process) of elementary finite cellular automata (i.e., $\{0, 1\}$ states, radius 1 and unidimensional) for which both states are quiescent (i.e. $(0, 0, 0) \mapsto 0$ and $(1, 1, 1) \mapsto 1$). It has been experimentally shown in previous works that introducing asynchronism in the global function of a cellular automata was perturbing its behavior, but as far as we know, only few theoretical work exists on the subject. The cellular automata we consider live on a ring of size n and asynchronism is introduced as follows: at each time step one cell is selected uniformly at random and the transition is made on this cell while the others stay in the same state. Among the sixty-four cellular automata belonging to the class we consider, we show that nine of them diverge almost surely on all non-trivial configurations while the fifty five other converge almost surely to a random fixed point. We show that the exact convergence time of these fifty five automata can only take the following values: either 0, $\Theta(n \ln n)$, $\Theta(n^2)$, $\Theta(n^3)$, or $\Theta(n2^n)$. Furthermore, the global behavior of each of these cellular automata is fully determined by reading its code.

Keywords: decentralized algorithm; asynchronous cellular automata; stochastic process; convergence predictions; distributed algorithms

Joint work of: Fatès, Nazim; Morvan, Michel; Schabanel, Nicolas; Thierry, Éric

"Fast" algorithms for Max Cut, Max 2-Sat, and other Max 2-CSPs

Gregory Sorkin (IBM TJ Watson Research Center, USA)

By Max 2-CSP we denote the set of maximization problems with arbitrary cost functions acting on pairs of binary variables; this class includes Max Cut, Max Dicut, Max 2-Sat, Max independent set, and weighted and minimization versions of all these problems. In this talk we show two results.

First, a Max 2-CSP with m clauses may be solved exactly in linear space and in time $O(m+n2^{m/5})$. The algorithm and analysis are simple, with the constant 5 given by the solution to a tiny linear program. A more sophisticated algorithm and analysis reduces $m/5$ to $19m/100$, improving on the previous special cases of $m/5$ for Max 2-Sat and $m/4$ for Max Cut.

Second, we show that a Max Cut of a $G(n, 1/n)$ random graph can be found by essentially the same algorithm in expected linear time. In fact, for any $\lambda > 0$ (possibly depending on n), any Max 2-CSP whose underlying graph is $G(n, c/n)$, $c \leq 1 + \lambda n^{-1/3}$, may be solved in expected time $O(n) \exp(O(1 + \lambda^3))$, and we conjecture that this is best possible. The key to the analysis is to obtain tail bounds on the "excess" of a component of a random graph by means of branching processes and random walks.

Joint work of: Sorkin, Gregory B.; Scott, Alexander D.

Approximation Techniques for Utilitarian Mechanism Design

Berthold Vöcking (RWTH Aachen, D)

We present approximation schemes for NP-hard mechanism design problems.

These algorithms need to satisfy certain monotonicity properties to ensure truthfulness. Since most of the known approximation techniques fail wrt. this property, we study alternative techniques.

Our main contribution is a quite general method to transform a pseudopolynomial time algorithm into a monotone FPTAS. This can be applied to various problems like, e.g. , knapsack, constrained shortest paths, or job scheduling with deadlines.

For example, the monotone FPTAS for the knapsack problem gives a very efficient, truthful mechanism for single-minded multi-unit auctions.

Joint work of: Vöcking, Berthold; Briest, Patrick; Krysta, Piotr

Balanced Boolean functions that can be evaluated so that every bit is unlikely to be read

David B. Wilson (Microsoft Research - Seattle, USA)

A Boolean function of n bits is balanced if it takes the value 1 with probability 1/2. We exhibit a balanced Boolean function with a randomized evaluation procedure (with probability 0 of making a mistake) so that on uniformly random inputs, no input bit is read with probability more than $\Theta(n^{-1/2}\sqrt{\log n})$.

We give a balanced monotone Boolean function for which the corresponding probability is $\Theta(n^{-1/3}\log n)$. We then show that for any randomized algorithm for evaluating a balanced Boolean function, when the input bits are uniformly random, there is some input bit that is read with probability at least $\Theta(n^{-1/2})$. For balanced monotone Boolean functions, there is some input bit that is read with probability at least $\Theta(n^{-1/3})$.

Joint work of: Wilson, David B.; Benjamini, Itai; Schramm, Oded

Approximate Distance Oracles and Spanners with sublinear surplus

Uri Zwick (Tel Aviv University, IL)

Several years ago, Thorup and Zwick obtained the following result:

Let $G = (V, E)$ be an undirected *weighted* graph with $|V| = n$ and $|E| = m$. Let $k \geq 1$ be an integer. Then, $G = (V, E)$ can be preprocessed in $O(kmn^{1/k})$ expected time, constructing a data structure of size $O(kn^{1+1/k})$, such that any subsequent *distance query* can be answered, approximately, in $O(k)$ time.

The approximate distance returned of stretch at most $2k-1$, i.e., it is at most $2k-1$ times the actual distance, and it is never too small. A girth conjecture of Erdős implies that $\Omega(n^{1+1/k})$ space is needed in the worst case for any stretch strictly smaller than $2k+1$. The space requirement of our algorithm is, hence, essentially optimal.

We now show that the techniques used to construct approximate distance oracles mentioned above for *weighted* graphs, can be used to obtain very simple constructions of spanners with *sublinear* error terms for *unweighted* graphs. These constructions extend, improve and simplify results of Elkin, Elkin and Peleg, and Bollobás, Coppersmith and Elkin.

More specifically, we show that for any integer $k > 1$, any undirected and unweighted graph $G = (V, E)$ on n vertices has a subgraph $G' = (V, E')$ with $O(kn^{1+1/k})$ edges such that for any two vertices $u, v \in V$, if $d_G(u, v) = d$, then $d_{G'}(u, v) = d + O(d^{1-1/(k-1)})$. (Here, $d_G(u, v)$ is the distance from u to v in G.) We also show that there is a *weighted* graph $G'' = (V, E'')$ with $O(kn^{1+1/(2^{k+1}-1)})$ edges such that for every $u, v \in V$, if $d_G(u, v) = d$, then $d \leq d_{G''}(u, v) = d + O(d^{1-1/(k-1)})$. The interesting feature of these new spanners is that the *relative* error decreases with the distance.

Joint work of: Thorup, Mikkel; Zwick, Uri