

08201 – Abstract Collection
Design and Analysis of Randomized and
Approximation Algorithms
— Dagstuhl Seminar —

Martin Dyer¹, Mark Jerrum² and Marek Karpinski³

¹ University of Leeds, GB
`dyer@scs.leeds.ac.uk`

² Queen Mary College - London, GB
`m.jerrum@qmul.ac.uk`

³ University of Bonn, DE
`marek@cs.uni-bonn.de`

Abstract. From 11.05.08 to 16.05.08, the Dagstuhl Seminar 08201 “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 work, and ongoing work and open problems were discussed. Abstracts of the presentations which were 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 paper are provided, if available.

Keywords. Randomized Algorithms, Approximation Algorithms, Optimization Problems, Measurement Problems, Approximation Complexity, Algorithmic Game Theory, Internet, Decentralized Networks, Network Design

08201 Overview – Design and Analysis of Randomized and Approximation Algorithms

The workshop 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. This included all sorts of completely new algorithmic questions that lie on the interface of several different areas. Here, some new broadly applicable techniques have emerged recently for designing efficient approximation algorithms for various optimization and measurement problems.

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

The 30 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 held in a very pleasant and stimulating atmosphere. Thanks to everyone who made it a very interesting and enjoyable event.

Martin Dyer
Mark Jerrum
Marek Karpinski

Acknowledgement. We thank Annette Beyer, Angelika Mueller-von Brochowski and Heike Clemens for their continuous support and help in organizing this workshop.

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

Dagstuhl Seminar 08201

Design and Analysis of Randomized and Approximation Algorithms

Monday, May 12th, 2008

09:00 - 09:10 Opening

Chair: Marek Karpinski

09:10 - 09:40 Alan Frieze (CMU - Pittsburgh)

Finding a Maximum Matching in a Sparse Random Graph in $O(n)$ Expected Time

09:40 - 10:10 Alexander Barvinok (University of Michigan)

Counting contingency tables and integer flows

10:10 - 10:40 Alex D. Scott (Oxford University)

Triangles in random graphs

10:40 - 11:00 Coffee break

Chair: Mark Jerrum

11:00 - 11:30 Michael Langberg (The Open Univ. of Israel - Raanan)

Two subgraph maximization problems

11:30 - 12:00 Zoya Svitkina (Dartmouth College - Hanover)

Sampling-based algorithms for submodular approximation

12:15 Lunch break

Chair: Martin Dyer

15:00 - 15:30 Catherine Greenhill (Univ. of New South Wales)

The cycle structure of two rows in a random latin square

15:30 - 16:00 Colin Cooper (King's College - London)

Multiple random walks on random regular graphs

16:00 - 16:30 Coffee break

18:00 Dinner

Dagstuhl Seminar 08201

Design and Analysis of Randomized and Approximation Algorithms

Tuesday, May 13th, 2008

Chair: Alan Frieze

- 09:00 - 09:30 Mark Jerrum (Queen Mary College - London)
An Approximation Trichotomy for Boolean #CSP
- 09:30 - 10:00 Leslie Ann Goldberg (University of Liverpool)
A complexity dichotomy for partition functions with mixed signs
- 10:00 - 10:30 Dimitris Achlioptas (Univ. California - Santa Cruz)
Algorithmic Phase Transitions in Random Constraint Satisfaction Problems
- 10:30 - 11:00 Coffee break

Chair: Michael Paterson

- 11:00 - 11:30 Christoph Dürr (CNRS, Ecole Polytechnique - Palaiseau)
Nash equilibria in Voronoi games on graphs
- 11:30 - 12:00 Ingo Wegener (Technische Universität Dortmund)
Tight Bounds for Blind Search on the Integers
- 12:15 Lunch break

Chair: Dorit Hochbaum

- 15:00 - 15:30 Thomas Hayes (Toyota Technological Institute)
The Forgiving Tree
- 15:30 - 16:00 Markus Jalsenius (University of Liverpool)
A systematic scan for 7-colourings of the grid
- 16:00 - 16:30 Coffee break

18:00 Dinner

Dagstuhl Seminar 08201

Design and Analysis of Randomized and Approximation Algorithms

Wednesday, May 14th, 2008

Chair: Moni Naor

09:00 - 09:30 Uriel Feige (Weizmann Inst. - Rehovot)
Structural Approximations: A Framework for Analyzing and Designing Heuristics

09:30 - 10:00 Bernd Gärtner (ETH Zürich)
Rational Unique Sink Orientations

10:00 - 10:30 Alantha L. Newman (RWTH Aachen)
Linear Equations mod p

10:30 - 11:00 Coffee break

Chair: Alexander Barvinok

11:00 - 11:30 Dorit S. Hochbaum (Univ. California - Berkeley)
Cut-based algorithms for image segmentation

11:30 - 12:00 Andrzej Lingas (Lund University)
Efficient approximation algorithms for shortest cycles in undirected graphs

12:15 Lunch break

13:30 - 17:30 Excursion

18:00 Dinner

19:30 Open Problem Session

Chair: Uriel Feige

Short talks given by:

Dimitris Achlioptas

Piotr Berman

Alan Frieze

Thomas Hayes

Dorith Hochbaum

Marek Karpinski

Christian Sohler

Zoya Svitkina

Dagstuhl Seminar 08201

Design and Analysis of Randomized and Approximation Algorithms

Thursday, May 15th, 2008

- Chair: Uriel Feige
- 09:00 - 09:30 Marek Karpinski (University of Bonn)
Sample Complexity of Nondense MAX-CUT
- 09:30 - 10:00 Moni Naor (Weizmann Inst. - Rehovot)
An Optimally Fair Coin Toss: Cleves Bound is Tight
- 10:00 - 10:30 Konstantinos Panagiotou (ETH Zürich)
Maximum Cuts and Other Extremal Subgraphs of Random Graphs
- 10:30 - 11:00 Coffee break
- Chair: Peter Gritzmann
- 11:00 - 11:30 Christian Sohler (University of Bonn)
Clustering for Metric and Non-Metric Distance Measures
- 11:30 - 12:00 Leen Stougie (TU Eindhoven)
Data gathering in wireless communication networks
- 12:15 Lunch break
- Chair: Artur Czumaj
- 15:00 - 15:30 Ola Svensson (IDSIA - Lugano)
(Acyclic) Job Shops are Hard to Approximate
- 15:30 - 16:00 Markus Bläser (Saarland University)
Randomness optimal identity tests of blackbox polynomials
- 16:00 - 16:30 Coffee break
- 18:00 Dinner

Dagstuhl Seminar 08201**Design and Analysis of Randomized and Approximation Algorithms****Friday, May 16th, 2008**

Chair: Moni Naor

09:00 - 09:30 Martin E. Dyer (University of Leeds)
*Randomly colouring random graphs*09:30 - 10:00 Magnus Bordewich (University of Durham)
*Path coupling without contraction*10:00 - 10:30 Amin Coja-Oghlan (University of Edinburgh)
Quasi-random graphs: eigenvalues and discrepancy

10:30 - 11:00 Coffee break

Chair: Leslie Ann Goldberg

11:00 - 11:30 Peter Gritzmann (TU München)
Optimal Wire Ordering and Spacing in Low Power Semiconductor Design

END OF WORKSHOP

12:15 Lunch

Finding a Maximum Matching in a Sparse Random Graph in $O(n)$ Expected Time

Alan M. Frieze (CMU - Pittsburgh, USA)

We present a linear expected time algorithm for finding maximum cardinality matchings in sparse random graphs. This is optimal and improves on previous results by a logarithmic factor.

Joint work with: Chebolu, Prasad; Melsted, Pall

Counting Contingency Tables and Integer Flows

Alexander Barvinok (University of Michigan, USA)

This is a report on the joint work with Alex Samorodnitsky, Zur Luria, and Alexander Yong. Given positive integer vectors $R = (r_1, \dots, r_m)$ and $C = (c_1, \dots, c_n)$ with $r_1 + \dots + r_m = c_1 + \dots + c_n = N$, we are interested in the number of $m \times n$ non-negative integer matrices (contingency tables) with the row sums R and column sums C . We present a randomized approximation algorithm which computes the number within the prescribed relative error $\epsilon > 0$. The complexity of the algorithm is $(1/\epsilon)^{O(1)} N^{O(\ln N)}$ for the class of “smooth margins” (R, C) . To define smooth margins, we consider the function $g(X) = \sum_{ij} (x_{ij} + 1) \ln(x_{ij} + 1) - x_{ij} \ln x_{ij}$ defined on positive $m \times n$ matrices $X = (x_{ij})$. The function g is strictly concave and hence attains its unique maximum X^* on the transportation polytope of non-negative $m \times n$ matrices with row sums R and column sums C . We say that the margins are smooth if the entries of $X^* = (x_{ij}^*)$ are not too large: $x_{ij}^* \leq \alpha N/mn$ for some constant $\alpha \geq 1$. The class of smooth margins includes, for example, the case where the ratios between the largest and the smallest row sums as well as between the largest and the smallest column sums are strictly smaller than the golden ratio $(1 + \sqrt{5})/2 \approx 1.618$.

Triangles in Random Graphs

Alex D. Scott (Oxford University, UK)

Let X be the number of triangles in a random graph $G(n, p)$. Loeb, Matousek and Pangrac showed that X is close to uniformly distributed modulo q when $p = 1/2$ and $q = O(\log n)$ is prime. We extend this result, showing that if $0 < p < 1$ is fixed, and $q = o(n^2/\log n)$ then X is asymptotically uniform. Furthermore, we discuss the implications of our methods for the distribution of X .

Two Subgraph Maximization Problems

Michael Langberg (Open Univ. of Israel - Raanan, IL)

In this talk I will discuss two natural maximization problems that have not been studied in the past. The first problem is the *Maximum Subgraph Homomorphism* problem - given two graphs G and H find a maximal edge subgraph of G that is H -colorable. This is a general problem with includes Max-Cut (when H is a single edge) and Max- k -Cut (when H is a k -clique). For MGH, I will present tight upper and lower bounds on the approximation ratio. Our *lower bounds* (hardness results) are not standard as they are based on average case complexity assumptions.

The second problem I will discuss is the Max- g -Girth problem - given a graph G find a maximum edge subgraph of G with girth at least g . Here, I also present both upper and lower bounds on the approximation ratio. However, the bounds presented are far from being tight, even for the case $g = 5$.

Joint work with: Kortsarz, Guy; Nutov, Zeev; Rabani, Yuval; Swamy, Chaitanya

Sampling-based algorithms for submodular approximation

Zoya Svitkina (Dartmouth College - Hanover, USA)

We introduce several generalizations of classical computer science problems, which involve a general submodular function in their objectives, in place of much simpler functions in the objectives of their classical counterparts. The problems include submodular knapsack, which generalizes the knapsack problem, and submodular load balancing, which generalizes load balancing or minimum-makespan scheduling. In these two problems, the size of a collection of items, instead of being just a sum of their individual sizes, is now a submodular function. Another two new problems are submodular sparsest cut and submodular balanced cut, which generalize their respective graph cut problems. Here, a general submodular function replaces the graph cut function, which itself is a well-known special case of a submodular function.

We establish upper and lower bounds for the approximability of these problems. Surprisingly, these factors are quite high. Whereas the corresponding classical problems are approximable to either constant or logarithmic factors, the guarantees that we prove for most of our algorithms are of the order of $\sqrt{(n/\log n)}$. We show that this is the inherent difficulty of these problems by proving matching (or, in some cases, almost matching) lower bounds.

Joint work with: Fleischer, Lisa

The cycle structure of two rows in a random latin square

Catherine Greenhill (Univ. of New South Wales)

Very little is known about the structure of a random latin square. We use switchings to investigate the cycle structure of the permutation defined by the first two rows of a random latin square. This permutation is a derangement, and we conjecture that the distribution of derangements formed by the first two rows of a latin square is asymptotically uniform (over the set of all derangements of a given size). While we fall short of proving this, we can show that the probability of seeing a given "typical" cycle structure is within a small polynomial multiplicative factor of the corresponding probability for random derangements.

Joint work with: Cavenagh, Nicholas; Wanless, Jan

Multiple Particle Random Walks in Random r -Regular Graphs

Colin Cooper (King's College London)

We study problems of oblivious and interactive particles making random walks on r -regular graphs. For non-interactive particles we prove that the speedup in covertime of the graph is exactly k , the number of particles. For interacting particles we consider various classical problems such as coalescing and annihilating particles in the vertex model. In these cases we prove the expected time to complete the process is $2\theta_r n$, $2 \ln \theta_r n$ and $2\theta_r n$ respectively, where $\theta_r = (r-1)/(r-2)$ and n is the number of vertices of the graph.

Joint work with: Frieze, Alan; Redzik, Tomasz

An Approximation Trichotomy for Boolean #CSP

Mark Jerrum (Queen Mary College - London, UK)

This talk examines the computational complexity of approximating the number of solutions to a Boolean constraint satisfaction problem (CSP). It extends a line of investigation started in a classical paper of Schaefer on the complexity of the decision problem for CSPs, and continued by Creignou and Hermann, who addressed exact counting. We find that the class of Boolean CSPs exhibits a trichotomy. Depending on the set of allowed relations, the CSP may be polynomial-time solvable (even exactly); or the number of solutions may be as hard to approximate as the number of accepting computations of a non-deterministic Turing machine. But there is a third possibility: approximating the number of solutions may be complete for a certain logically defined complexity class, and hence equivalent in complexity to a number of natural approximate counting problems, of which independent sets in a bipartite graph is an example.

A Complexity Dichotomy for Partition Functions with Mixed Signs

Leslie Ann Goldberg (University of Liverpool, UK)

Partition functions (also known as *homomorphism functions*) form a rich family of graph invariants including the number of k -colourings of a graph, the number of independent sets of a graph, and also the partition functions of certain spin-system models from statistical physics, such as the Ising model.

Building on earlier work by Dyer and Greenhill and Bulatov and Grohe, we classify the computational complexity of partition functions. Our main result is a dichotomy theorem stating that every partition function is either computable in polynomial time or $\#P$ -complete. Partition functions are described by symmetric matrices with real entries, and we prove that it is decidable in polynomial time in terms of the matrix whether a given partition function is in polynomial time or $\#P$ -complete.

While in general it is very complicated to give an explicit algebraic or combinatorial description of the tractable cases, for partition functions described by Hadamard matrices — these turn out to be central in our proofs — we obtain a simple algebraic tractability criterion, which says that the tractable cases are those *representable* by a quadratic polynomial over a 2-element finite field.

Joint work with: Grohe, Martin; Jerrum, Mark; Thurley, Marc

Algorithmic Phase Transitions in Random Constraint Satisfaction Problems

Dimitris Achlioptas (Univ. California - Santa Cruz, USA)

For many random Constraint Satisfaction Problems by now we have a good understanding of the largest constraint density for which solutions exist. At the same time, though, all known polynomial-time algorithms for these problems stop finding solutions at much smaller densities. We study this phenomenon by examining how the different sets of solutions evolve as constraints are added. We prove in a precise mathematical sense that, for each problem studied, the barrier faced by algorithms corresponds to a phase transition in that problem's solution-space geometry. Roughly speaking, at some problem-specific critical density, the set of solutions shatters and goes from being a single giant ball to exponentially many, well-separated, tiny pieces. All known polynomial-time algorithms work in the ball regime, but stop as soon as the shattering occurs. Besides shedding some light on the failure of algorithms studied until now, our results also provide new constructions of one-way functions.

Nash Equilibria in Voronoi Games on Graphs

Christoph Dürr (CNRS, Ecole Polytechnique - Palaiseau, F)

We study a game where every player is to choose a vertex (facility) in a given undirected graph. All vertices (customers) are then assigned to closest facilities and a player's payoff is the number of customers assigned to it. We show that deciding the existence of a Nash equilibrium for a given graph is NP-hard which to our knowledge is the first result of this kind for a zero-sum game. We also introduce a new measure, the social cost discrepancy, defined as the ratio of the costs between the worst and the best Nash equilibria. We show that the social cost discrepancy in our game is $\Omega(\sqrt{(n/k)})$ and $O(\sqrt{(kn)})$, where n is the number of vertices and k the number of players.

Tight Bounds for Blind Search on the Integers

Ingo Wegener (Technische Universität Dortmund, D)

We analyze a simple random process in which a token is moved in the interval $A = \{0, n\}$: Fix a probability distribution μ over $\{1, \dots, n\}$. Initially, the token is placed in a random position in A . In round t , a random value d is chosen according to μ . If the token is in position $a \geq d$, then it is moved to position $a - d$. Otherwise it stays put. Let T be the number of rounds until the token reaches position 0. We show tight bounds for the expectation of T for the optimal distribution μ , i.e., we show that $\min_{\mu} \{E\mu(T)\} = \Theta((\log n)^2)$. For the proof, a novel potential function argument is introduced. The research is motivated by the problem of approximating the minimum of a continuous function over $[0, 1]$ with a 'blind' optimization strategy.

The Forgiving Tree

Thomas Hayes (Toyota Technological Institute, USA)

We present two distributed algorithms for self-healing in networks that are reconfigurable (such as peer-to-peer networks) in the sense that they can change their topology during an attack. Self-healing seeks to maintain connectivity and possibly other useful properties in the face of repeated attacks by an adversary, that in our model, is assumed to be omniscient.

The first algorithm DASH is an algorithm that adds edges only among neighbors of deleted nodes (i.e. locality-aware) and provably maintains connectivity and limits the degree increase of any node. Our second algorithm ForgivingTree allows only $O(1)$ degree increase and also limits the diameter increase of the network. It, however, may need to use nodes other than the neighbors of the deleted node for reconstruction. Our approach is orthogonal and complementary to traditional topology-based approaches to defending against attack.

Joint work with: Trehan, Amitabh; Rustagi, Navin; Saida, Jared

A Systematic Scan for 7-Colourings of the Grid

Markus Jalsenius (University of Liverpool, UK)

This talk examines the mixing time of a systematic scan Markov chain for sampling from the uniform distribution on proper 7-colourings of a finite subregion G of the infinite square lattice, the grid. A systematic scan Markov chain cycles through finite-size subsets of vertices of G in a deterministic order and updates the colours assigned to the vertices of each subset. The systematic scan Markov chain that we present cycles through subsets consisting of 2×2 -subgrids. The colours assigned to the vertices of the 2×2 -subgrids are updated using a procedure known as heat bath. We give a computer-assisted proof that this systematic scan Markov chain mixes in $O(\log n)$ scans, where n is the number of vertices of G . We make use of a heuristic to compute required couplings of distributions on colourings of 2×2 -subgrids. This is the first time the mixing time of a systematic scan Markov chain on the grid has been shown to mix rapidly for less than 8 colours. We also give partial results that underline the challenges of proving rapid mixing of a systematic scan Markov chain for sampling 6-colourings of the grid by considering 2×3 - and 3×3 -subgrids.

Joint work with: Pedersen, Kasper

Structural Approximations: A Framework for Analyzing and Designing Heuristics

Uriel Feige (Weizmann Institut - Rehovot, IL)

We introduce a new framework for designing and analyzing algorithms. Our framework applies best to problems that are inapproximable according to the standard worst-case analysis. We circumvent such negative results by designing guarantees for classes of instances, parameterized according to properties of the optimal solution. We also make sure that our parameterized approximation, called the structural approximation, is the best possible. We show how to apply our framework to problems with additive and submodular objective functions such as the set buying and maximum facility location problems. We consider two types of algorithms for these problems. For greedy algorithms, our framework provides a justification for preferring a certain natural greedy rule over some alternative greedy rules that have been used in similar contexts. For LP-based algorithms, we show that the natural LP relaxation for these problems is not optimal in our framework. We design a new LP relaxation and show that this LP relaxation coupled with a new randomized rounding technique is optimal in our framework.

Rational Unique Sink Orientations

Bernd Gärtner (ETH Zürich, CH)

An orientation of the n -cube graph is called a unique sink orientation (USO) if every nonempty cube face has a unique sink. The algorithmic challenge is to find the global sink, using as few vertex evaluations as possible (a vertex evaluation is a call to an oracle that reveals for the given vertex the orientations of the incident edges).

In this talk we introduce a class of USO that are generated by finite state transducers and therefore have succinct representations. We discuss the question how fast the well-known Random-Facet algorithm performs on these and provide experimental evidence that the performance is superpolynomial on some family of USO generated by a transducer with just 3 states.

Relaxed Linear Equations mod p

Alantha Newman (Dimacs, Rutgers University)

We consider an optimization problem that is related to the classic Linear Equations mod p Problem. Suppose we are given a set of equations of the form $\{x_j - x_i \equiv c_{ij} \pmod{p}\}$. Our goal is to satisfy these equations as much as possible. Specifically, suppose each x_i is assigned a value $x_j - x_i \equiv c_{ij} \pm y_{ij} \pmod{p}$. Then we want to maximize the value $1 - 2y_{ij}/p$ over all equations.

We discuss two different approaches to this problem, both SDP-based. The first has a guarantee of 0.637. The second, we conjecture to be close to the "optimal" value of 0.878.

Cut-Based Algorithms for Image Segmentation

Dorit S. Hochbaum (University California - Berkeley, USA)

The Markov Random Field (MRF) is a well-known model for image segmentation. The idea is to achieve uniform color areas (clusters of pixels) in an image by assigning a separation penalty function to different color assignments of adjacent pixels that is increasing with increasing difference. A second component of the penalty function is to penalize the deviation between the given color (the prior) and the assigned color. It is shown that when the functions are convex the problem is solvable in polynomial time as a convex dual of the min cost network flow problem. For linear separation functions (such as absolute value) the problem is solvable with a best possible complexity as parametric minimum cut.

Additional problems known in clustering and in imaging are:

- (1) To find a subset S in a graph with similarity weights on the edges, that minimizes the ratio of the minimum cut $c(S, \bar{S})$ divided by the similarity within the set S $c(S, S)$.

- (2) To minimize $c_{S \subset V} \frac{c(S, \bar{S})}{|S|}$ ("ratio cut").

Both these problems are shown to be solvable in (strongly) polynomial time as a min-cut on a certain related graph. Both these problems were previously believed to be NP-hard.

Efficient Approximation Algorithms for Shortest Cycles in Undirected Graphs

Andrzej Lingas (Lund University, S)

We describe a simple combinatorial approximation algorithm for finding a shortest (simple) cycle in an undirected graph. For an undirected graph G of unknown girth k , our algorithm returns with high probability a cycle of length at most $2k$ for even k and $2k + 2$ for odd k , in time $O(n^{3/2} \sqrt{\log n})$. Thus, in general, it yields a $8/3$ approximation. We study also the problem of finding a simple cycle of minimum total weight in an undirected graph with nonnegative edge weights. We present a simple combinatorial 2-approximation algorithm for a minimum weight (simple) cycle in an undirected graph with nonnegative integer edge weights in the range $\{1, 2, \dots, M\}$, where $M > n$. This algorithm runs in time $O(n^2 \log n \log M)$.

Sample Complexity of Nondense MAX-CUT

Marek Karpinski (University of Bonn, D)

We present some new results on the sample complexity and approximability of nondense instances of MAX-CUT and MAX-CSP. Among other things, we prove existence of relative approximation schemes for general instances of MAX-CUT and MAX-CSP with running times depending explicitly only on their density.

Joint work with: Fernandez de la Vega, W.

An Optimally Fair Coin Toss: Cleves Bound is Tight

Moni Naor (Weizmann Institut - Rehovot, IL)

We address one of the classical and long-standing open problems in the foundations of cryptography: the bias of coin-flipping protocols. Coin-flipping protocols allow mutually distrustful parties to generate a common unbiased random bit, guaranteeing that even if one of the parties is malicious, it cannot significantly bias the output of the honest party.

A classical result by Cleve [STOC '86] showed that for any two-party coin-flipping protocol with r rounds, there exists an efficient adversary that can bias the output of the honest party by $\Omega(1/r)$. However, the best previously known

protocol only guarantees $O(1/\sqrt{r})$ bias, and the question of whether Cleve's bound was tight has remained open for more than twenty years.

In this work we establish the optimal tradeoff between the round complexity and the maximal bias of two-party coin-flipping protocols. Under standard assumptions, we show that Cleve's lower bound is tight: we construct an r -round protocol with bias $O(1/r)$.

Joint work with: Moran, Tal; Segev, Gil

Maximum Cuts and Other Extremal Subgraphs of Random Graphs

Konstantinos Panagiotou (ETH Zürich, CH)

Let K_l denote the complete graph on l vertices. We prove that there is a constant $c = c(l) > 0$, such that whenever $p \geq n^{-c}$, with probability tending to 1 when n goes to infinity, every maximum K_l -free subgraph of the binomial random graph $G_{n,p}$ is $(l-1)$ -partite.

The proof is based on a result of independent interest: we show, for instance, that the maximum cut of almost all graphs with M edges, where $M \gg n$, is nearly unique. More precisely, given a maximum cut C of $G_{n,M}$, we can obtain all maximum cuts by moving at most $O(\sqrt{n3/M})$ vertices between the parts of C .

Clustering for Metric and Non-Metric Distance Measures

Christian Sohler (University of Bonn, D)

We study a generalization of the k -median problem with respect to arbitrary distance measure D . Given a finite point set P , our goal is to find a set C of size k such that the sum of error $D(P, C) = \sum_{p \in P} \min_{c \in C} D(p, c)$ is minimized.

We obtain a $(1 + \epsilon)$ -approximation algorithm with $O\left(n2^{(k/\epsilon)^{O(1)}}\right)$ running time for this problem, if the 1-median problem can be approximated within a factor of $1 + \epsilon$ by taking a random sample of constant size and solving the 1-median problem on the sample exactly. Using this characterization, we obtain the first linear time $(1 + \epsilon)$ -approximation algorithm for the k -median problem for the Kullback-Leibler divergence and a number of other distance measures.

Data Gathering in Wireless Communication Networks

Leen Stougie (TU Eindhoven, NL)

The Wireless Gathering Problem is to find an interference-free schedule for data gathering in a wireless network in minimum time. We present a 4-approximate polynomial-time on-line algorithm for this NP-hard problem. We show that no shortest path following algorithm can have approximation ratio better than 4.

(Acyclic) Job Shops are Hard to Approximate

Ola Svensson (IDSIA - Lugano, I)

Currently, the best approximation algorithms for job shops have slightly worse than logarithmic performance guarantee, and the only known inapproximability result says that it is NP-hard to approximate acyclic job shops within a factor less than $5/4$.

We narrow this gap by providing new inapproximability results. More specifically, we show that the (acyclic) job shop problem cannot be approximated within ratio $O(\log^{1-\epsilon} lb)$, unless NP has (randomized) quasi-polynomial algorithms, and where lb denotes a trivial lower bound on the optimal value. This almost matches the best known results for acyclic job shops, since an $O(\log^{1+\epsilon} lb)$ -approximate solution can be obtained in polynomial time.

Recently, a PTAS was given for the job shop problem, where the number of machines and the number of operations per job are assumed to be constant. We show that both these restrictions are necessary to obtain a PTAS, even in the preemptive case, i.e., when jobs are allowed to be interrupted.

Joint work with: Mastrolilli, Monaldo

Randomness Optimal Identity Tests of Blackbox Polynomials

Markus Bläser (Saarland University, D)

We give the first asymptotically optimal blackbox identity test for multivariate polynomials with degrees d_1, \dots, d_n in the variables X_1, \dots, X_n . Our algorithms use $(1 + o(1))OPT$ random bits for dense polynomials and for broad classes of sparse polynomials. The runtime of our algorithm is logarithmic in the degree of the input polynomial. The basis for our result results are several novel constructions of hitting set generators.

Randomly Colouring Random Graphs

Martin E. Dyer (University of Leeds, UK)

We consider the problem of generating a colouring of the random graph $\mathbb{G}_{n,p}$ uniformly at random using a natural Markov chain algorithm: the Glauber dynamics. We assume that there are $\beta\Delta$ colours available, where Δ is the maximum degree of the graph, and we wish to determine the least $\beta = \beta(p)$ such that the distribution is close to uniform in $O(n \log n)$ steps of the chain. This problem has been previously studied for $\mathbb{G}_{n,p}$ in cases where np is relatively small. Here we consider the *dense* cases, where $np \in [om \ln n, n]$ and $om = om(n) \rightarrow \infty$. Our methods are closely tailored to the random graph setting, but we obtain

considerably better bounds on $\beta(p)$ than can be achieved using more general techniques.

Joint work with: Frieze, Alan

Path Coupling without Contraction

Magnus Bordewich (University of Durham, UK)

Path coupling is a useful technique for simplifying the analysis of a coupling of a Markov chain. Rather than defining and analysing the coupling on every pair in $\Omega \times \Omega$, where Ω is the state space of the Markov chain, analysis is done on a smaller set $S \subseteq \Omega \times \Omega$. If the coefficient of contraction β is strictly less than one, no further analysis is needed in order to show rapid mixing. However, if $\beta = 1$ then analysis (of the variance) is still required for all pairs in $\Omega \times \Omega$. We present a new approach which shows rapid mixing in the case $\beta = 1$ with a further condition which only needs to be checked for pairs in S , greatly simplifying the work involved.

Joint work with: Dyer, Martin

Quasi-Random Graphs: Eigenvalues and Discrepancy

Amin Coja-Oghlan (University of Edinburgh, UK)

A graph $G = (V, E)$ with density $p = \#E/\binom{\#V}{2}$ satisfies $\text{DISC}(\epsilon)$ if

$$\forall S \subset V \quad |2e(S) - \#S^2 p| < 2\epsilon \#E,$$

where $e(S)$ is the number of edges inside of S . We show that $\text{DISC}(\gamma\epsilon^2)$ implies the following:

There is a set $W \subset V$, $\#W \leq (1 - \epsilon)\#V$ such that the matrix $M = (M_{vw})_{v,w \in W}$ with

$$M_{vw} = P - \begin{cases} 1 & \text{if } \{v, w\} \in E \\ 0 & \text{otherwise} \end{cases}$$

satisfies $\|M\| < \epsilon np$.

Here γ is a constant.

Joint work with: Alon, Noga.; Han, Hiep; Kang, Miun; Rödl, Vojtech; Schacht, Mathias

Optimal Wire Ordering and Spacing in Low Power Semiconductor Design

Peter Gritzmann (TU München, D)

A key issue for high integration circuit design in the semiconductor industry are power constraints that stem from the need for heat removal and reliability or battery lifetime limitations. As the power consumption depends heavily on the capacitances between adjacent wires, determining the optimal ordering and spacing of parallel wires is an important issue in the design of low power chips.

As it turns out, optimal wire spacing is a convex optimization problem while the combined optimal wire spacing and ordering is a special class of the minimal Hamilton-path problem. While the latter is NP-hard in general, the present paper provides an $O(N \log N)$ algorithm that solves the underlying coupled ordering and spacing problem for N parallel wires to optimality.

Joint work with: Ritter, Michael; Zuber, Paul