

Seminar on Combinatorics, Games and Optimisation in 2017/18

Seminars are listed in reverse chronological order, most recent first

Thursday 26 July - [Vasilis Gkatzelis](#) (Drexel)

Deferred-Acceptance Auctions: Worst-Case Approximation Guarantees

Deferred-acceptance auctions are mechanisms whose allocation rule needs to use an adaptive reverse greedy algorithm. Economists recently introduced these auctions and proved that they satisfy remarkable incentive guarantees that make them very practical. However, we have a very limited understanding of the extent to which such reverse greedy algorithms can provide desirable worst-case approximation guarantees. Computer scientists have analyzed several forward greedy algorithms in the past but, as we show, reverse greedy algorithms are much more complicated (e.g., they do not even guarantee maximality). In this talk we first study the limitations of this class of algorithms, and then we design new approximation algorithms from this class, which induce novel deferred-acceptance auctions.

Thursday 19 July - [Ahmad Abdi](#) (University of Waterloo)

Testing idealness of 0,1 matrices

A 0,1 matrix is *ideal* if the corresponding set covering polyhedron is integral. In a surprising turn of events, Ding, Feng and Zang (2008) showed that testing idealness of a 0,1 matrix is co-NP-complete. The reason for this hardness result is that finding certain non-ideal substructures is NP-complete.

In yet another surprising turn of events, we show that the blockers of the same non-ideal substructures, however, can be detected in A in polynomial time. In this talk, I will survey the current state of affairs, and sketch the proof of our result.

This is joint work with Gerard Cornuejols and Dabeen Lee.

Thursday 28 June - [Tom Lidbetter](#) (Rutgers Business School)

Burning spiders, forests and more: graph burning as a model of social contagion.

Graph burning is a recently introduced model for the spread of memes and

contagion in social networks. The burning process on a graph starts with all vertices unburned at time $t=0$. At each time step $t=1,2,\dots$, one new unburned vertex is set on fire. In each new round, all vertices that are adjacent to some burning vertex catch on fire. The process ends when all vertices are on fire. The burning number of a graph is the least number of rounds needed to burn all the vertices. It is conjectured that any graph has burning number at most the ceiling of \sqrt{n} , but the proof is elusive. We make modest progress here by settling the conjecture for spider graphs (trees with exactly one vertex of degree more than 2).

This is joint work with Anthony Bonato.

Thursday 7 June - Alex Fink (QMUL)

The Tutte polynomial via lattice point enumeration

I will explain how to recover the Tutte polynomial of a matroid from an Ehrhart-style polynomial which counts lattice points in Minkowski sums of simplices and its base polytope. The key ingredient is a polyhedral interpretation of activity; along the way, this will give us a regular subdivision whose cells naturally encode Dawson's activity partition.

I will also talk about its generalisation to polymatroids: in this setting, finding a bivariate activity invariant was a question of Tamás Kálmán, who constructed the univariate activity invariant in his work on enumerating spanning trees of hypergraphs.

This is joint work with Amanda Cameron (MPI Leipzig).

Thursday 24 May - Istvan Tomon (EPFL)

On the size of k -cross-free families

Two subsets A, B of a finite ground set X are said to be crossing, if none of the four sets $A \cap B$, $A \setminus B$, $B \setminus A$ and $X \setminus (A \cup B)$ are empty. A family of subsets of X is k -cross-free, if it does not contain k pairwise crossing elements. The notion of k -cross-free families was first introduced by Karzanov in the context of multicommodity flow problems, but these families also appear in combinatorial geometry, and in evolutionary biology.

It was conjectured by Karzanov and Lomonosov forty years ago that if a family F of subsets of the n -element ground set X does not contain k pairwise crossing elements, then $|F| = O(kn)$. For $k = 2$ and 3 , the conjecture is true, but for larger values of k the best known upper bound, due to Lomonosov, is $|F| = O(kn \log n)$. We improve this bound by showing that $|F| = O_k(n \log^{**} n)$ holds, where \log^{**} denotes the iterated logarithm function.

This is joint work with Andrey Kupavskii and János Pach.

Thursday 17 May - Robert Simon (LSE)

Paradoxical decompositions and finitary rules

We colour every point x of a probability space X according to the colours of a finite list x_1, x_2, \dots, x_k of points such that each of the x_i , as a function of x , is a measure preserving transformation. We ask two questions about a colouring rule: (1) does there exist a finitely additive extension of the probability measure for which the x_i remain measure preserving and also a colouring obeying the rule almost everywhere that is measurable with respect to this extension?, and (2) does there exist any colouring obeying the rule almost everywhere? If the answer to the first question is no and to the second question yes, we say that the colouring rule is paradoxical. A paradoxical colouring rule not only allows for a paradoxical partition of the space, it requires one.

Thursday 26 April - Georg Loho (EPFL)

Monomial tropical cones for multicriteria optimization

We introduce a special class of tropical cones which we call 'monomial tropical cones'. They arise as a helpful tool in the description of discrete multicriteria optimization problems. After an introduction to tropical convexity with an emphasis on these particular tropical cones, we explain the algorithmic implications. We finish with connections to commutative algebra.

Wednesday 18 April - Fabrizio Grandoni (IDSIA Lugano)

Approximating Geometric Knapsack via L-packings

In the 2-dimensional geometric knapsack problem (2DK) we are given a set of n axis-aligned rectangular items, each one with an associated profit, and an axis-aligned square knapsack. The goal is to find a (non-overlapping) packing of a maximum profit subset of items inside the knapsack (without rotating items). The best-known polynomial-time approximation factor for this problem (even just in the cardinality case) is $2 + \epsilon$ [Jansen and Zhang, SODA 2004]. In this work we break the 2 approximation barrier, achieving a polynomial-time $17/9 + \epsilon < 1.89$ approximation.

Essentially all prior work on 2DK approximation packs items inside a constant number of rectangular containers, where items inside each container are packed using a simple greedy strategy. We deviate for the first time from this setting: we show that there exists a large profit solution where items are packed inside a constant number of containers plus one L-shaped region at the boundary of the knapsack, which contains items that are high and narrow or wide and thin.

As a second contribution, we present a PTAS for packing items in this L-shaped region. The previous best approximation factor for this subproblem was $2+\epsilon$ (obtained via a trivial reduction to 1-dimensional knapsack by considering tall or wide items only).

We achieve better approximation factors for the cardinality case and/or when we are allowed to rotate items by 90 degrees. The previous best approximation factor for all such variants was $2 + \epsilon$ as well [Jansen and Zhang, SODA 2004].

This is joint work with: Waldo Galvez, Sandy Heydrich, Salvatore Ingala, Arindam Khan, and Andreas Wiese.

Thursday 19 April - [Eilon Solan](#) (Tel Aviv University)

Optimal Repeated Inspection

We study a discounted repeated inspection game with two agents and one principal. Both agents may profit by violating certain rules, while the principal can inspect at most one agent in each period, inflicting a punishment on an agent who is caught violating the rules. The goal of the principal is to minimize the discounted number of violations, and she has a Stackelberg leader advantage. We characterize the principal's optimal inspection strategy. It turns out that the optimal history-dependent inspection strategy lowers the discounted loss due to violations to a small fraction of the performance of inspection strategies that were studied in the literature.

Joint work with Chang Zhao.

Wednesday 28 March - [Jefferson Huang](#) (Cornell)

Dynamic Scheduling and Maintenance of a Deteriorating Server

Motivated by a quality control problem in semiconductor manufacturing, we consider a stochastic scheduling problem in the context of a multi-class queue with a single server whose service capacity deteriorates randomly over time. We show that the system may be unstable under a natural extension of the $c\mu$ -rule, and provide a sufficient condition for this rule to be optimal. We also consider the problem of jointly deciding whether to perform service or preventive maintenance, for which we provide insights into the structure of optimal policies and heuristics.

Wednesday 21 March - [Jugal Garg](#) (UIUC)

Fisher Markets and Nash Social Welfare

Fisher market equilibrium is a powerful solution concept that has found several surprising applications even in non-market settings which do not involve any

exchange of money but only require the remarkable fairness and efficiency properties of equilibria, e.g., scheduling, auctions, mechanism design, fair division, etc. A very recent new application is approximation algorithms for maximizing the Nash social welfare (NSW) when allocating a set of indivisible items to a set of agents.

In this talk, I will start with the Fisher market model and its connection with the NSW problem. Then, I will show how to design a constant-factor approximation algorithm for maximizing the NSW when agents have budget-additive valuations. Budget-additive valuations represent an important class of submodular functions. They have attracted a lot of research interest in recent years due to many interesting applications.

This is based on a joint work with Martin Hoefer and Kurt Mehlhorn.

Thursday 8 March - [Robert Simon](#) (LSE)

Games of Incomplete Information and Myopic Equilibria

What happens if the payoffs of an infinitely repeated game are the sums of some payoffs on the first n stages and of some other payoffs from the undiscounted game? If it is a game of incomplete information on one side, there will be a Nash equilibrium. To prove this, a new equilibrium concept is introduced, that of myopic equilibria. With myopic equilibria an agent can become locked into severely suboptimal strategies by the expectations of others.

Joint work with S. Spiez, H. Torunczyk.

Wednesday 7 March - [Sergei Chubanov](#) (Siegen)

A polynomial scaling algorithm for linear programming

In this talk we will consider a polynomial algorithm for linear programming based on a scaling technique. As well as the ellipsoid method, this algorithm is applicable to linear problems given by separation oracles. We will also discuss some numerical results concerning machine learning applications and the perspectives of this algorithm in infinite-dimensional linear optimization.

Thursday 1 March - [Andrey Kupavskii](#) (Birmingham)

Lower Bounds for Searching Robots, some Faulty

Suppose we are sending out k robots from 0 to search the real line at a constant speed (with turns) to find a target at an unknown location; f of the robots are faulty, meaning that they fail to report the target although visiting its location. The goal is to find the target in time at most $\lambda |d|$, if the target is located at d , $|d| > 1$, for

λ as small as possible. In this work, we find a tight lower bound for λ . Joint work with Emo Welzl.

Thursday 22 February - Katherine Staden (Oxford)

The minimum number of triangles in a graph of given order and size

A famous theorem of Mantel from 1907 states that every n -vertex graph with more than $n^2/4$ edges contains at least one triangle. In the 50s, Erdős asked for a quantitative version of this statement: for every n and e , how many triangles must an n -vertex e -edge graph contain? This question has received a great deal of attention, and a long series of partial results culminated in an asymptotic solution by Razborov, extended to larger cliques by Nikiforov and Reiher. Until recently, an exact solution was only known for a small range of edge densities, due to Lovász and Simonovits. In this talk, I will discuss the history of the problem and some new work which gives an exact solution for almost the entire range of edge densities. This is joint work with Hong Liu and Oleg Pikhurko.

Wednesday 21 February - Peyton Young (LSE and Oxford)

The Speed of Innovation Diffusion in Social Networks

New technologies typically gain a foothold through the actions of a few innovators, and then diffuse more rapidly as more and more people come into contact with prior adopters.

Much of the prior literature focuses on the rate of diffusion as a function of the topology of a given network structure. Here we derive "topology-free" bounds on the expected waiting time until a given fraction of the population has adopted the innovation.

The bounds depend on the payoff gain from using the innovation instead of the status quo, and on the noisiness of the players' response functions, but they do not depend on the network structure per se. In particular, the bounds hold for directed and undirected networks of arbitrary size whose structure may be evolving over time.

Thursday 15 February - Fatemeh Mohammadi (Bristol)

Chip-firing Games: A Combinatorial and Geometric Perspective

The chip firing game is a solitaire game on graphs which has a very beautiful and rich mathematical structure, and it appears in several fields of mathematics and statistics. It starts from a configuration of dollars (chips) on the vertices of a fixed graph: in each step of a chip-firing game we may choose a vertex to lend one dollar to each of its neighbours, or to borrow one dollar from each of its neighbours. The goal of the game is to get all the vertices out of debt (with non-negative integers) by

a sequence of legal moves. For any finite graph, there are only finitely many critical configurations whose vertices are not in debt and no subset of vertices can fire without going into debt. I will talk about critical configurations, their corresponding geometric objects and their applications in system reliability theory.

Thursday 8 February - [Trine Tørnøe Platz](#) (Copenhagen Business School)

On totally balanced, submodular and PMAS-admissible weighted minimum colouring games

We introduce the weighted minimum colouring (WMC) games, which is a class of cooperative combinatorial optimization games. A graph $G = (N, E)$ and a positive integer weight vector w that assigns a weight to each vertex in N induce a WMC game. Our aim is to characterize classes of graphs that induce WMC games with specific properties for either any choice of weight vector or for at least one weight vector. A graph G is said to be globally (respectively, locally) WMC totally balanced, submodular, or PMAS-admissible, if for all positive integer weight vectors w (respectively, for at least one positive integer weight vector w), the corresponding WMC game is totally balanced, submodular, or admits a PMAS (population monotonic allocation scheme). We show that a graph G is globally WMC totally balanced if and only if it is perfect, and that any graph G is locally WMC totally balanced. Furthermore, we show that G is globally (respectively, locally) WMC submodular if and only if it is complete r -partite (respectively, $(2K_2, P_4)$ -free). Finally, we show that G is globally PMAS-admissible if and only if it is $(2K_2, P_4)$ -free.

Joint work with Herbert Hamers, Nayat Horozoglu, and Henk Norde.

Wednesday 7 February - [Tamas Kiraly](#) (Eotvos University)

Tight $\sqrt{2}$ -approximation for Linear 3-Cut

In the linear 3-cut problem, the input is a node-weighted directed graph and three specified terminal nodes s, r, t , and the goal is to find a minimum weight subset of non-terminal nodes whose removal ensures that s cannot reach r and t , and r cannot reach t . This problem is approximation-equivalent to the following arborescence blocking problem: given a node-weighted directed graph with a specified root node r , remove a minimum weight subset of non-root nodes such that the remaining digraph has no in-arborescence and no out-arborescence rooted at r .

Linear 3-cut contains undirected 3-way node cut as a special case, and can be reduced to directed 2-way cut. Under the Unique Games Conjecture, the best approximation ratios of these problems are $4/3$ and 2 , respectively. We show that the linear 3-cut problem has a $\sqrt{2}$ -approximation algorithm and this is tight under UGC. The proof involves showing that, somewhat surprisingly, the integrality gap of

the natural LP-relaxation is $\sqrt{2}$. Joint work with Kristóf Bérczi, Karthekeyan Chandrasekaran, and Vivek Madan.

Thursday 1 February - Igal Milchtaich (Bar-Ilan University)

Polyequilibrium

Polyequilibrium is a generalization of Nash equilibrium that is applicable to any strategic game, whether finite or otherwise, and to dynamic games, with perfect or imperfect information. It differs from equilibrium in specifying strategies that players do not choose and by requiring an after-the-fact justification for the exclusion of these strategies rather than the retainment of the non-excluded ones.

Specifically, for each excluded strategy of each player there must be a non-excluded one that responds to every profile of non-excluded strategies of the other players at least as well as the first strategy does. A polyequilibrium's description of the outcome of the game may be more or less specific, depending on the number and the identities of the retained, non-excluded strategy profiles. A particular result (e.g., Pareto efficiency of the payoffs) is said to hold in a polyequilibrium if it holds for all non-excluded profiles. Such a result does not necessarily hold in any Nash equilibrium in the game. In this sense, the generalization proposed in this work extends the set of justifiable predictions concerning a game's results.

Wednesday 17 January - David C. Parkes (Paulson School of Engineering and Applied Sciences, Harvard University)

Optimal Economic Design through Deep Learning

*****This seminar was part of the [Social and Economic Data Science Seminars Series](#)*****

Designing an auction that maximizes expected revenue is an intricate task. Despite major efforts, only the single-item case is fully understood. We explore the use of tools from deep learning on this topic. The design objective that we adopt is revenue optimal, dominant-strategy incentive compatible auctions. For a baseline, we show that multi-layer neural networks can learn almost-optimal auctions for a variety of settings for which there are analytical solutions, and even without leveraging characterization results. We also show that deep learning can be used to derive auctions for poorly understood problems, including settings with multiple items and budget constraints. Our research also demonstrates that the deep learning framework is quite general, being applicable to other problems of optimal economic design.

Joint work with Paul Duetting (LSE), Zhe Feng (Harvard University), and Harikrishna Narasimhan (Harvard University). Working paper: <https://arxiv.org/abs/1706.03459>

Thursday 11 January - Johannes Carmesin (Cambridge)

Embedding simply connected 2-complexes in 3-space

We characterise the embeddability of simply connected 2-dimensional simplicial complexes in 3-space in a way analogous to Kuratowski's characterisation of graph planarity, by excluded minors. This answers questions of Lovász, Pardon and Wagner

**Wednesday 29 November - in combination with SEDs Data Science Seminar series
Elisa Celis and Nisheeth Vishnoi (EPFL)**

Controlling Bias in Artificial Intelligence

Bias is an increasingly observed phenomenon in the world of artificial intelligence (AI) and machine learning: From gender bias in online search to racial bias in court bail pleas to biases in worldviews depicted in personalized newsfeeds. How are societal biases creeping into the seemingly "objective" world of computers and programs? At the core, what is powering today's AI are algorithms for fundamental computational problems such as classification, data summarization, and online learning. Such algorithms have traditionally been designed with the goal of maximizing some notion of "utility" and identifying or controlling bias in their output has not been a consideration. In this talk, Nisheeth and Elisa will explain the emergence of bias in algorithmic decision making and present the first steps towards developing a systematic framework to control biases in several of the aforementioned problems. This leads to new algorithms that have the ability to control and alleviate bias, often without a significant compromise to the utility that the current algorithms obtain.

Wednesday 22 November - Justin Ward (Queen Mary)

Improved Approximation for K-Means in Arbitrary Dimension

In the general K-Means problem in which we are given n input points in a Euclidean space and seek to find k "center" points in the space so that the sum of the squared distances of each input point to its nearest center is minimized. While there have been several recent results for the special case in which the Euclidean space has some bounded dimension d , the best known approximation in arbitrary Euclidean spaces has remained $(9+\epsilon)$ since 2002. In this talk I will present a new algorithm that achieves a 6.36-approximation for this problem, as well as an improved 2.64 approximation for the Euclidean k -median problem. The algorithm is based on a new Lagrangian multiplier preserving primal dual approach. This talk is based on joint work with Sara Ahmadian, Ashkan Norouzi-Fard, and Ola Svensson.

Thursday 16 November - Tim Roughgarden (Stanford & LSE)

Distribution-Free Models of Social and Information Networks

The mathematical study of social and information networks has historically centered around generative models for such networks (preferential attachment, the

Chung-Lu random graph model, Kronecker graphs, etc.). This talk proposes distribution-free models of social and information networks --- novel classes of graphs that capture all plausible such networks. Our models are motivated by triadic closure, the property that vertices with one or more mutual neighbors tend to also be neighbors; this is one of the most universal signatures of social networks. We prove structural results on the clustering properties of such graphs, and give algorithmic applications to clustering and clique-finding problems.

Includes joint work with Jacob Fox, Rishi Gupta, C. Seshadhri, Fan Wei, and Nicole Wein.

9 November - Christoph Koch (Warwick)

How to re-establish a brand: bootstrap percolation in social networks

We consider a classical and well-studied graph process called $\text{\emph{bootstrap percolation}}$. While bootstrap percolation originally arose in the context of disordered magnetic systems, it has become an important model in the description and analysis of dynamics in social networks. Initially some vertices in a graph are $\text{\emph{activated}}$, and subsequently each vertex with at least ≥ 2 active neighbours is activated as well. Once a vertex is active, it remains so forever.

On a variety of random graph models it is known that even (reasonably) sparse initial activations are likely to cause an activation of most vertices in the graph. However, for other models this $\text{\emph{critical phenomenon}}$ does not occur. We investigate this dichotomy in detail and determine almost tight criteria on the degree sequence of random graphs for both the existence and non-existence of this critical phenomenon.

These results are joint work with M. Kang (TU Graz), T. Makai (TU Graz), and N. Fountoulakis (Birmingham).

2 November - Liana Yepremyan (Oxford)

Rainbow matchings in properly edge-colored multigraphs

Aharoni and Berger conjectured that in any bipartite multigraph that is properly edge-coloured by n colours with at least $n + 1$ edges of each colour there must be a matching that uses each colour exactly once (such a matching is called rainbow).

This conjecture recently have been proved asymptotically by Pokrovskiy. In this talk I will consider the same question without the bipartiteness assumption. It turns out that in any multigraph with bounded edge multiplicities that is properly edge-coloured by n colours with at least $n + o(n)$ edges of each colour there must be a matching of size $n - O(1)$ that uses each colour at most once. This is joint work with Peter Keevash.

1 November - Peter Csikvari (Budapest)

Schrijver's theorem on the number of perfect matchings and its variants

In this talk we will sketch a new proof of Schrijver's theorem. This theorem asserts that a d -regular bipartite graph on $2n$ vertices has at least $C_d n$ perfect matchings, where $C_d = (d-1)^{(d-1)} / d^{(d-2)}$. (I will explain where the constant comes from.) The new proof uses ideas from graph limit theory, and relies on the work of Heilmann and Lieb concerning the matching polynomial. Then we will survey several further applications of the method.

25 October - Péter Pál Pach (Warwick/Technical University Budapest)

On some recent applications of the polynomial method

In this talk we will look at a new variant of the polynomial method which was first used to prove that sets avoiding 3-term arithmetic progressions in groups like \mathbb{Z}_4^n and \mathbb{F}_q^n are exponentially small (compared to the size of the group). Since then many interesting applications of this method were shown, for instance, the solution of the Erdős-Szemerédi sunflower conjecture, tight bound for Green's arithmetic triangle removal lemma and growth rate of tri-colored sumfree sets. Finally, I will also mention some open problems.

26 October - He Sun (Edinburgh)

Heat kernels in graphs: A journey from random walks to geometry, and back

Heat kernels are one of the most fundamental concepts in physics and mathematics. In physics, the heat kernel is a fundamental solution of the heat equation and connects the Laplacian operator to the rate of heat dissipation. In spectral geometry, many fundamental techniques are based on heat kernels. In finite Markov chain theory, heat kernels correspond to continuous-time random walks and constitute one of the most powerful techniques in estimating the mixing time. In this talk, we will briefly discuss this line of research and its relation to heat kernels in graphs. In particular, we will see how heat kernels are used to design the first nearly-linear time algorithm for finding clusters in real-world graphs. Some interesting open questions will be addressed in the end.

11 October - Haris Aziz (University of New South Wales, Sydney)

A discrete and bounded envy-free cake cutting protocol for any number of agents

We consider the well-studied cake cutting problem in which the goal is to find an envy-free allocation. The problem has received attention in computer science, mathematics, and economics. It had been an open problem whether there exists a discrete and bounded envy-free protocol. In this talk I will discuss our discrete and bounded envy-free protocol that resolved the problem.

(Based on joint work with Simon Mackenzie)

12 October - Michelle Delcourt (Birmingham)

On the List Coloring Version of Reed's Conjecture

Reed conjectured in 1998 that the chromatic number of a graph should be at most the average of the clique number (a trivial lower bound) and maximum degree plus one (a trivial upper bound); in support of this conjecture, Reed proved that the chromatic number is at most some nontrivial convex combination of these two quantities. King and Reed later showed that a fraction of roughly $1/130000$ away from the upper bound holds. Motivated by a paper by Bruhn and Joos, last year Bonamy, Perrett, and Postle proved for large enough maximum degree, a fraction of $1/26$ away from the upper bound holds, a significant step towards the conjectured value of $1/2$. Using new techniques, we show that the list-coloring version holds; for large enough maximum degree, a fraction of $1/13$ suffices for list chromatic number. This result implies that $1/13$ suffices for ordinary chromatic number as well. This is joint work with Luke Postle.

6 October - Pablo Moscato (University of Newcastle, Australia)

"We have it all wrong"... so what are you doing to change practice?

Along with many other researchers, I share the view that a systematically coherent research program, in both theory and applications of algorithms, is definitely needed to accelerate innovation in computing. We routinely design computational approaches and engage in healthy competitions where the performance of our methods is tested... but what if "We have it all wrong"? What if we need a paradigmatic change in our practice for the development and design of computational methods? We may need to enrich our practice with a new approach.

In fact, John N. Hooker already alerted the computing and mathematical community more than 20 years ago [Hooker, 1995; Journal of Heuristics]: "Competitive testing tells us which algorithm is faster but not why." Hooker argued for a more scientific approach and he proposed the use of 'controlled experimentation'. This is common in empirical sciences. "Based on one's insights into an algorithm", he said, "one may

expect good performance to depend on a certain problem characteristic". Then "design a controlled experiment that checks how the presence or absence of this characteristic affects performance" and, finally, "build an exploratory mathematical model that captures the insight [...] and deduce from its precise consequences that can be put to the test". In this talk, I will address how a new thinking is needed for the development of our field. I will have an with emphasis in our success on both speeding up solutions for the traveling salesman problem as well as our success to create very hard instances for the world's fastest solver.

4 October - Martin Hoefer (Frankfurt)

Algorithms for Nash Social Welfare and Markets with Satiation

Nash social welfare is an interesting objective with many appealing properties for item allocation in markets. The algorithmic problem of optimizing or approximating this objective, however, is only poorly understood. In this talk, I discuss the first constant-factor approximation algorithm when agents have budget-additive valuation functions. Budget-additive functions are an important class of submodular functions, and they attracted significant interest in approximation algorithms and algorithmic game theory in recent years.

Our algorithm relies on rounding an approximate equilibrium in a linear Fisher market where sellers have earning limits (upper bounds on the amount of money they want to earn) and buyers have utility limits (upper bounds on the amount of utility they want to achieve). In contrast to markets with either earning or utility limits, these markets have not been studied before, and they have fundamentally different properties. We design an FPTAS to compute an approximate equilibrium, a result that may be of independent interest.

29 September - Jens Vygen (Bonn)

Approaching $3/2$ for the s-t-path TSP

The s-t-path TSP is a variant of the traveling salesman problem in which the endpoints of the tour are given and distinct. The integrality ratio of the natural linear programming relaxation is believed to be $3/2$, but all approximation algorithms known so far have worse performance ratio. We show that there is a polynomial-time algorithm with approximation guarantee $3/2 + \epsilon$, for any fixed $\epsilon > 0$.

It is well known that Wolsey's analysis of Christofides' algorithm also works for the s-t-path TSP except for the narrow cuts (in which the LP solution has value less than two). A fixed optimum tour has either a single edge in a narrow cut (then call the edge and the cut lonely) or at least three (then call the cut busy). Our algorithm "guesses" (by dynamic programming) lonely cuts and edges. Then we partition the instance into smaller instances and strengthen the LP, requiring value at least three

for busy cuts. By setting up a k -stage recursive dynamic program, we can compute a spanning tree (V,S) and an LP solution y such that $1/2 + O(2^{-k})y$ is in the T -join polyhedron, where T is the set of vertices whose degree in S has the wrong parity.

This is joint work with Vera Traub.