Department of Mathematics

Seminar on Combinatorics, Games and Optimisation in 2018

Seminars are listed in reverse chronological order, most recent first

---

Thursday 13 December - Vera Traub (Hausdorff Centre for Mathematics)

Beating the integrality ratio for s-t-tours in graphs
Among various variants of the traveling salesman problem, the s-t-path graph TSP has the special feature that we know the exact integrality ratio, 3/2, and an approximation algorithm matching this ratio. In this work, we go below this threshold: we devise a polynomial-time algorithm for the s-t-path graph TSP with approximation ratio 1.497.
Our algorithm can be viewed as a refinement of the 3/2-approximation algorithm by Sebő and Vygen [2014], but we introduce several completely new techniques. These include a new type of ear-decomposition, an enhanced ear induction that reveals a novel connection to matroid union, a stronger lower bound, and a reduction of general instances to instances in which s and t have small distance (which works for general metrics).

This is joint work with Jens Vygen.

Wednesday 12 December - Noam Nisan (Hebrew University of Jerusalem)

The Communication Complexity of Cake-Cutting
This talk concerns the well-studied model of "cake-cutting" for studying questions regarding notions of fair division of resources. We focus on discrete versions rather than using infinite-precision real values, specifically, focusing on the communication complexity. Using general discrete simulations of classical infinite-precision protocols (Robertson-Webb and moving-knife), we roughly partition the various fair-allocation problems into 3 classes: "easy" (constant number of rounds of logarithmic many bits), "medium" (poly-log total communication), and "hard".
Our main technical result concerns two of the "medium" problems (perfect

allocation for 2 players and equitable allocation for any number of players) which we prove are not in the "easy" class. Our main open problem is to separate the "hard" from the "medium" classes.

Joint work with Simina Brânzei

Thursday 6 December - Andrés Cristi (Universidad de Chile)

A Near Optimal Mechanism for Energy Aware Scheduling
Consider a cloud server, where clients can submit jobs for processing. The quality of service that each agent receives is given by a private non-decreasing function of the completion time of her job. The server has to process the jobs and charge each agent while trying to optimize the social cost that is defined as the energy expenditure plus the sum of the values of the internal cost functions. The server operator would like to design a mechanism in order to optimize this objective, which ideally is computationally tractable, charges the users “fairly” and the induced game has an equilibrium. We present a mechanism that combines the aforementioned properties with a constant Price of Anarchy. An interesting feature of our mechanism is that it is indirect: each user needs only to declare an upper bound on the completion time of her job, and not the cost function.

This is joint work with Antonios Antoniadis.

Friday 30 November - Benny Sudakov (ETH)

Subgraph statistics
Consider integers k,ℓ such that $0\le \ell \le \binom{k}2$. Given a large graph G, what is the fraction of k-vertex subsets of G which span exactly ℓ edges? When G is empty or complete, and ℓ is zero or $\binom k 2$, this fraction can be exactly 1. On the other hand if ℓ is not one these extreme values, then by Ramsey's theorem, this fraction is strictly smaller than 1.
The systematic study of the above question was recently initiated by Alon, Hefetz, Krivelevich and Tyomkyn who proposed several natural conjectures. In this talk we discuss a theorem which proves one of their conjectures and implies an 0asymptotic version of another. We also make some first steps towards analogous question for hypergraphs. Our proofs involve some Ramsey-type arguments, and a number of different probabilistic tools, such as polynomial anticoncentration inequalities and hypercontractivity. Joint work with M. Kwan and T. Tran

Wednesday 28 November - Guillem Perarnau Llobet (Birmingham)

Efficient sampling of random colorings

A well-known conjecture in computer science and statistical physics is that Glauber dynamics on the set of k-colorings of a graph G on n vertices with maximum degree \Delta is rapidly mixing for k \ge \Delta+2. In 1999, Vigoda showed rapid mixing of flip dynamics with certain flip parameters on the set of proper k-colorings for k > (11/6)\Delta, implying rapid mixing for Glauber dynamics. In this paper, we obtain the first improvement beyond the (11/6)\Delta barrier for general graphs by showing rapid mixing for k > (11/6 - \eta)\Delta for some positive constant \eta. The key to our proof is combining path coupling with a new kind of metric that incorporates a count of the extremal configurations of the chain. Additionally, our results extend to list coloring, a widely studied generalization of coloring. Combined, these results answer two open questions from Frieze and Vigoda's 2007 survey paper on Glauber dynamics for colorings. (Joint work with Michelle Delcourt and Luke Postle)

Wednesday 21 November - Iannis Caragiannis (University of Patras)

Fairness in allocation problems

In this talk, we will present variations of the problem of allocating indivisible goods to agents with additive valuations for the goods. We will define basic fairness concepts such as proportionality and envy-freeness and will discuss their properties. Next, we will focus on allocations of maximum Nash welfare, and we will explain how they guarantee approximate notions of envy-freeness. Finally, we will define new fairness concepts that are related to the level of awareness of the agents for the allocation and to their social relations. We will present examples and many open problems. No advanced or special mathematical background is needed to follow the talk.

Thursday 15 November - Tibor Szabó (Freie Universität Berlin)

Optimality of the uniform random strategy

The concept of biased Maker-Breaker games, introduced by Chvátal and Erd\H os, is a central topic in the field of positional games, with deep connections to the theory of random structures. For any given hypergraph H the main questions is to determine the smallest bias q(H) that allows Breaker to force that Maker ends up with an independent set of H. Here we prove matching general winning criteria for Maker and Breaker when the game hypergraph satisfies a couple of natural

‘container-type’ regularity conditions about the degree of subsets of its vertices. This will enable us to derive a hypergraph generalization of the H-building games, studied for graphs by Bednarska and Luczak.
Furthermore, we investigate the biased version of generalizations of the van der Waerden games introduced by Beck. We refer to these generalizations as Rado games and determine their threshold bias up to constant factors by applying our general criteria. We find it quite remarkable that a purely game theoretic deterministic approach provides the right order of magnitude for such a wide variety of hypergraphs, when the generalizations to hypergraphs in the analogous setup of sparse random discrete structures are usually quite challenging.

Wednesday 14 November - Matthias Englert (Warwick)

Polylogarithmic Guarantees for Generalized Reordering Buffer Management
In the Generalized Reordering Buffer Management Problem (GRBM) a sequence of items located in a metric space arrive online, and have to be processed by a set of k servers moving within the space. In a single step the first b still unprocessed items from the sequence are accessible, and a scheduling strategy has to select an item and a server.
Then the chosen item is processed by moving the chosen server to its location. The goal is to process all items while minimizing the total distance travelled by the servers.This problem was introduced by Chan et al. (TCS'12) and has been subsequently studied in an online setting by Azar et al. (STACS'14). The problem is a natural generalization of two very well-studied problems: the k-server problem for b=1 and the Reordering Buffer Management Problem (RBM) for k=1. We consider the GRBM problem on a uniform metric in the online version.
We show how to obtain a competitive ratio of O(log k(log k+loglog b)) for this problem. This is a significant improvement in the dependency on b compared to the previous best bound of O(sqrt(b) log k), and is asymptotically optimal for constant k.

Thursday 8 November - Maryam Sharifzadeh (Warwick)

Number of maximal sum-free subsets of integers and abelian groups.
A set S is sum-free if it does not contain a triple x,y,z such that $x+y=z$ (note x,y may not necessarily be distinct). We say that $S\subseteq [n]$ is a maximal sum-free subset of $[n]$ if it is sum-free and it is not properly contained in another sum-free subset of $[n]$. Cameron and Erd\H{o}s asked whether the number of maximal sum-free subsets of $[n]$ is much smaller than the number of sum-free

sets. They also gave a lower bound of $2^{\lfloor n/4 \rfloor }$ for the number of maximal sum-free sets. Together with J\'ozsef Balogh, Hong Liu, and Andrew Treglown, we give an exact solution to the problem: For each $1\leq i \leq 4$, there is a constant C_i such that, given any $n\equiv i \mod 4$, $[n]$ contains $(C_i+o(1)) 2^{n/4}$ maximal sum-free sets. In this talk, I will outline the ideas and techniques used in the proof. I will also discuss some new results on the number of maximal sum-free sets in abelian groups, which are joint work with Hong Liu.

Wednesday 7 November - Brendan Murphy (Bristol)

Solymosi's conjecture for rich lines in general position
We give an explicit construction that disproves a conjecture of Solymosi on the number of lines in general position that contain a near maximal number of points of a Cartesian product set. We will explain how this problem is connected to the sum-product problem, and show why Solymosi's conjecture is plausible. The counter-example we give is motivated by related questions in group theory, and is rather different from the usual examples in arithmetic combinatorics.

Thursday 1 November - Kristof Berczi (Eotvos University)

Improving the integrality gap for multiway cut
In the multiway cut problem, we are given an undirected graph with non-negative edge weights and a collection of k terminal nodes, and the goal is to partition the node set of the graph into k non-empty parts each containing exactly one terminal so that the total weight of the edges crossing the partition is minimized. For arbitrary k, the best-known approximation factor is 1.2965 due to Sharma and Vondrák while the best known inapproximability factor is 1.2 due to Angelidakis, Makarychev and Manurangsi.
In this talk we show how to improve on the lower bound by constructing an integrality gap instance for the CKR relaxation. A technical challenge in improving the gap has been the lack of geometric tools to understand higher-dimensional simplices. We analyze the gap of the instance by viewing it as a convex combination of 2-dimensional instances and a uniform 3-dimensional instance. One of the byproducts from our proof technique is a generalization of a result on Sperner admissible labelings due to Mirzakhani and Vondrák.

Joint work with Karthakeyan Chanrdasekaran, Tamás Király and Vivek Madan.

Wednesday 31 October - Ágnes Cseh (IE CERS HAS)

The complexity of cake cutting with unequal share

An unceasing problem of our prevailing society is the fair division of goods. The problem of proportional cake cutting focuses on dividing a heterogeneous and divisible resource, the cake, among n players who value pieces according to their own measure function. The goal is to assign each player a not necessarily connected part of the cake that the player evaluates at least as much as her proportional share.

We investigate the problem of proportional division with unequal shares, where each player is entitled to receive a predetermined portion of the cake. Our main contribution is threefold. First we present a protocol for integer demands that delivers a proportional solution in fewer queries than all known algorithms. Then we show that our protocol is asymptotically the fastest possible by giving a matching lower bound. Finally, we turn to irrational demands and solve the proportional cake cutting problem by reducing it to the same problem with integer demands only. All results remain valid in a highly general cake cutting model, which can be of independent interest.

Joint work with Tamás Fleiner. The paper received the best paper award at SAGT 2018.

Thursday 25 October - Bruce Reed (McGill)

ω, Δ, and x

We consider three graph invariants and the relationship between them. The first is the size of the largest clique in a graph G, denoted $\omega(G)$. The second is the maximum degree of a vertex of G, denoted $\Delta(G)$. The third is the chromatic number of G, denoted $x(G)$.Trivially, $\omega(G) \leq x(G)$. Greedy colouring shows $x(G) \leq \Delta(G)+1$. So $\omega(G) \leq x(G) \leq \Delta(G)+1$. We discuss a conjecture which states that x lies in the lower half of this range. We also show that when x is sufficiently close to Δ it is determined by a small subgraph and can be computed in polynomial time.

Wednesday 24 October - Varun Kanade (University of Oxford)

Hierarchical clustering: Objectives & Algorithms

Hierarchical clustering is a recursive partitioning of a dataset into clusters at an increasingly finer granularity. Hierarchical clustering has mostly been studied in procedural terms, i.e., a hierarchical cluster tree is obtained using certain bottom-

up (agglomerative) or top-down (divisive) heuristics, rather than finding a hierarchical cluster tree as the solution obtained by optimizing some suitable objective. Dasgupta (2016) identified this as a reason why the theory of hierarchical clustering lagged behind that of flat clustering and proposed an objective function. In this talk, we will take an axiomatic approach to identifying suitable objective functions for hierarchical clustering. We will also describe a new random-graph model with a planted hierarchy. New and existing algorithms and their performance in theory and on some preliminary experiments will be discussed.

Based on joint work with Vincent Cohen-Addad, Claire Mathieu and Frederik Mallmann-Trenn.

Thursday 18 October - Leonidas Pitsoulis (Aristotle University of Thessaloniki)

Decomposition of signed-graphic matroids
In this seminar we will present a decomposition theory for the class of signed graphic matroids which are representable either in GF(2) or GF(4). The proposed decomposition differs from previous decomposition results on matroids that have appeared in the literature in the sense that it is not solely based on k-sums such as the decomposition of regular matroids, but also on an operation called star composition. A sketch of the resulting recognition algorithms as well as an excluded minor characterization of the building blocks of the aforementioned decomposition will also be presented.

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,..., 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 Tam\'as K\'alm\'an, 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 $x_1, x_2, \ldots, x_k$ 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 $2 + \varepsilon$ [Jansen and Zhang, SODA 2004]. In this work we break the 2 approximation barrier, achieving a polynomial-time $17/9 + \varepsilon < 1.89$ 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 $2+\varepsilon$ (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 $2 + \varepsilon$ 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 cμ-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 lambda as small as possible. In this work, we find a tight lower bound for lambda. 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 dept 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 Tornoe 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 √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 √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 √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