Department of Mathematics

Seminar and PhD Seminar on Combinatorics, Games and Optimisation in 2019

Seminars are listed in reverse chronological order, most recent first

---

Friday 13 December - Gal Kronenberg (University of Oxford)
Venue: 32L.B.09 from 12:00 - 13:00

Turán numbers of long cycles in random graphs
For a graph G on n vertices and a graph H, denote by ex(G,H) the maximal number of edges in an H-free subgraph of G. We consider a random graph G~G(n,p) where p=C/n, and study the typical value of ex(G,H), where H is a long cycle.
We determine the asymptotic value of ex(G,C_t), where G~G(n,p), p> C/n and A*log(n)< t< (1- Ɛ)n. The behaviour of ex(G,C_t) can depend substantially on the parity of t. In particular, our results match the classical result of Woodall on the Turán number of long cycles, and can be seen as its random version. In fact, our techniques apply in a more general sparse pseudo-random setting.
We also prove a robustness-type result, showing the likely existence of cycles of prescribed lengths in a random subgraph of a graph with a nearly optimal density.

Joint work with Michael Krivelevich and Adva Mond.

Thursday 12 December - Spyros Angelopoulos (CNRS and Sorbonne University)

Online Computation with Untrusted Advice
The advice model of online computation captures the setting in which an online algorithm is given some partial information concerning the request sequence. This paradigm allows to establish tradeoffs between the amount of this additional information and the performance of the online algorithm. However, unlike real life in which advice is a recommendation that we can choose to follow or to ignore based on trustworthiness, in the current advice model, the online algorithm typically treats it as infallible. This means that if the advice is corrupt or, worse, if it comes from a malicious source, the algorithm may perform poorly. In this work, we study online computation in a setting in which the advice is provided by an untrusted source. Our objective is to quantify the impact of untrusted advice so as to design and analyze robust online algorithms that are resilient and perform well even when the advice is generated in a malicious, adversarial manner. We show how the new paradigm can be applied to well-studied online problems such as ski rental, online

bidding, bin packing, and list update.

Joint work with Christoph Dürr, Shendan Jin, Shahin Kamali and Marc Renault.

Friday 6 December - Bento Natura (LSE)

A scaling-invariant algorithm for linear programming whose running time depends only on the constraint matrix
Following the breakthrough work of Tardos (Oper. Res. '86) in the bit-complexity model, Vavasis and Ye (Math. Prog. '96) gave the first exact algorithm for linear programming in the real model of computation with running time depending only on the constraint matrix. For solving a linear program (LP), Vavasis and Ye developed a primal-dual interior point method using a layered least squares step, and showed that O(n^3.5 * log chi) iterations suffice to solve linear programs exactly, where chi is a condition measure controlling the size of solutions to linear systems related to the constraint matrix.

Monteiro and Tsuchiya (SIAM J. Optim. '03), noting that the central path is invariant under rescalings of the columns of the constraint matrix and the objective function, asked whether there exists an LP algorithm depending instead on the measure chi*, defined as the minimum chi value achievable by a column rescaling of the constraint matrix, and gave strong evidence that this should be the case. We resolve this open question affirmatively.

We will illustrate the central ideas on how to develop a scaling-invariant algorithm for LP and how to find a near-optimal rescaling.

Thursday 5 December - Victor Verdugo (LSE)

Prophet Inequalities and Mechanism Design
In many situations finding the optimal revenue pricing policy requires to solve a hard optimisation problem. Posted price mechanism are simple and efficiently implementable. In this talk I'll show the connection between this type of mechanisms and optimal stopping rules for online selection problems, and how the guarantees from one problem to the other are preserved.

Thursday 28 November - László Végh (LSE)

An improved constant-factor approximation algorithm for the asymmetric travelling salesman problem
In this talk, I will speak about a recent result by Vera Traub and Jens Vygen. For the asymmetric variant of the classical travelling salesman problem, the first constant-factor approximation algorithm was given in our 2018 paper with Ola Svensson and Jakub Tarnawski. The result used a combination of combinatorial and linear programming tools, and reduced the problem to more structured special cases in a

number of reduction steps.

The approximation ratio in the first version of our paper was 5500, which we later improved to 506. The new paper by Traub and Vygen dramatically improves this to 22. Whereas they follow the same overall strategy as our previous paper, they manage to find substantial simplifications in the chain of reductions that not only improve the performance guarantee, but also make the result more accessible.

Friday 22 November - Nora Frankl (LSE)

On the number of discrete chains in the plane
Determining the maximum number of unit distances that can be spanned by n points in the plane is a difficult problem, which is wide open. The following more general question was recently considered by Eyvindur Ari Palsson, Steven Senger, and Adam Sheffer. For given distances t_1,...,t_k a (k+1)-tuple (p_1,...,p_{k+1}) is called a k-chain if ||x_i-x_{i+1}||=t_i for i=1,...,k. What is the maximum possible number of k-chains that can be spanned by a set of n points in the plane? Improving the result of Palsson, Senger and Sheffer, we determine this maximum up to a small error term (which, for k=1 mod 3 involves the maximum number of unit distances). We also consider some generalisations, and the analogous question in R^3. Joint work with Andrey Kupvaskii.

Thursday 21 November - Joonkyung Lee (Universität Hamburg)

Convex graphon parameters and graph norms
Sidorenko's conjecture states that the number of a bipartite graph H in a graph~G is asymptotically minimised when G is a quasirandom graph. A notorious example that this conjecture remains open is the case $H=K_{5,5}\setminus C_{10}$. It has been even unknown whether this graph possesses the weakly norming property, a strictly stronger property than satisfying the conjecture.
We take a step towards understanding the graph $K_{5,5}\setminus C_{10}$ by proving that it is not weakly norming. More generally, we show that `twisted' blow-ups of cycles, which include $K_{5,5}\setminus C_{10}$ and $C_6\square K_2$, are not weakly norming. This answers two questions of Hatami, who asked whether the two graphs are weakly norming. The method relies on analysing Hessian matrices defined by graph homomorphisms, by using the equivalence between the (weakly) norming property and convexity of graph homomorphism densities. We also prove that $K_{t,t}$ minus a perfect matching, proven to be weakly norming by Lov\'asz, is not norming for every~$t>3$. Joint work with Bjarne Sch\"ulke.

Wednesday 20 November - Franziska Eberle (Universität Bremen)

Commitment in online scheduling made easy
We study a fundamental online job admission problem where jobs with processing times and deadlines arrive online over time at their release dates, and the task is to determine a preemptive single-server schedule which maximizes the number of jobs that complete on time. To circumvent known impossibility results, we make a standard slackness assumption by which the feasible time window for scheduling a job is at least (1+epsilon) times its processing time, for some epsilon>0. We consider a variant of the online scheduling problem where the provider has to satisfy certain commitment requirements. These requirements arise, for example, in modern cloud-services, where customers do not want last-minute rejections of critical tasks and request an early-enough provider-side commitment to completing admitted jobs.

Our main contribution is an optimal algorithm for online job admission with commitment. When a provider must commit upon starting a job, our bound is O(1/epsilon). This is best possible as there is a lower bound of Omega(1/epsilon) for online admission even without commitment. If the commitment decisions must be made before a job's slack becomes less than a delta-fraction of its size, we prove a competitive ratio of O(epsilon/((epsilon - delta) delta)) for 0 <delta<epsilon. This result interpolates between commitment upon starting a job and commitment upon arrival. For the latter commitment model, it is known that no (randomized) online algorithm does admit any bounded competitive ratio.

Friday 15 November - Cosmin Pohoata (Caltech)

Sets without 4APs but with many 3APs
It is a classical theorem of Roth that every dense subset of $\left\{1,\ldots,N\right\}$ contains a nontrivial three-term arithmetic progression. Quantitatively, results of Sanders, Bloom, and Bloom-Sisask tell us that subsets of relative density at least $1/(\log N)^{1-\epsilon}$ already have this property. In this talk, we will discuss about some sets of N integers which unlike $\left\{1,\ldots,N\right\}$ do not contain nontrivial four-term arithmetic progressions, but which still have the property that all of their subsets of relative density at least $1/(\log N)^{1-\epsilon}$ must contain a three-term arithmetic progression. Perhaps a bit surprisingly, these sets turn out not to have as many three-term progressions as one might be inclined to guess, so we will also address the question of how many three-term progressions can a four-term progression free set may have. Finally, we will also discuss about some related results over $\mathbb{F}_{q}^n$. Based on joint works with Jacob Fox and Oliver Roche-Newton.

Thursday 14 November - Eliza Jablonska (University of Cracow)

On generically Haar-"small" sets in Abelian Polish groups
A subset A of an Abelian Polish group X is *Haar-null* (following Christensen) if there are a Borel set B covering A and a Borel probability measure m on X such that $m(x+B)=0$ for all x in X.
Dodos proves that for every Haar-null Borel subset A of X the set of all test measures

$$T(A):=\{m \in P(X)\ :\ m(x+A)=0 \text{ for all } x \in X\}$$

is dense, coanalytic, and either meagre or co-meagre in the space $P(X)$ of all probability Borel measures on X (with Lévy-Prokhorov metric).
Christensen's measure-theoretic notion has a topological analogue due to Darji: a subset A of X is *Haar-meagre* if there are a Borel set B covering A, a compact metric space K and a continuous function $f: K \to X$ such that

$$f^{-1}(B+x) \text{ is meagre in } K \text{ for all } x \in X$$

We prove an analogue of the Dodos theorem: for every Haar-meager Borel subset A, the set of all witness functions

$$W(A):=\{ f \in C(2^{\omega},X) :\ f^{-1}(x+A) \text{ is meagre in } 2^{\omega} \text{ for all } x \in X\}$$

is dense, coanalytic, and either meager or comeager in in the space $C(2^{\omega}, X)$ of all continuous functions $f: 2^{\omega} \to X$ with the supremum metric. The first step is to show that the compact metric space K in Darji's definition can always be replaced by the Cantor cube 2^{ω}.

Wednesday 13 November - Gyorgy Turan (UIC)

Interpretability in machine learning
In many applications of machine learning, learned models or their decisions need to be interpretable (or explainable, comprehensible). For example, `why was my credit application rejected?'. Neural networks, for example, are typically not interpretable, while decision trees are more so. We give a general overview of the topic, and discuss some recent projects.
The theoretical part is about efficient approximation of Bayesian networks with `interpretable' models.

Friday 8 November - Olaf Parczyk (LSE)

The size-Ramsey number of tight 3-uniform paths
Given a hypergraph H, the size-Ramsey number is the smallest integer m such that there exists a graph G with m edges with the property that in any colouring of the

edges of G with two colours there is a monochromatic copy of H. Extending on results for graphs we prove that the size Ramsey number of the 3-uniform tight path on n vertices is linear in n.

This is joint work with Jie Han, Yoshiharu Kohayakawa, and Guilherme Mota.

Thursday 7 November - Miquel Oliu-Barton (Université Paris-Dauphine)

A Solution for Stochastic Games
Shapley (1953) introduced stochastic games in order to model dynamic interactions in which the environment changes in response to the players' behavior, and proved that finite competitive stochastic games admit a Î»-discounted value for any discount rate Î». The case where Î» is close to zero is of particular interest, as it corresponds to an interaction in the long run, far from opportunistic behavior. Bewley and Kohlberg (1976) proved that the Î»-discounted values converge as Î» goes to zero. Building on this result, Mertens and Neyman (1981) proved that finite competitive stochastic games admit a value, and that the value coincides with the limit of the Î»-discounted values as Î» goes to zero. Finding a tractable formula for the value of finite competitive stochastic games was a major open problem for nearly 40 years. The present contribution resolves this problem.

Joint work with Luc Attia.

Friday 1 November - Anupam Gupta (CMU)

K-way-cuts in graphs
For an undirected graph, a k-way cut is a set of edges whose deletion breaks the graph into at least k pieces. How fast can we find a minimum-weight k-way cut? And how many minimum k-way cuts can a graph have? The two problems seem to be closely linked --- in 1996 Karger and Stein showed how to find a minimum k-way cut in time approximately n^{2k-2}, and also that the number of minimum k-way cuts is at most n^{2k-2}. Both these results are not known to be tight, except for the case of k=2, that of finding graph min-cuts.
In this talk, we report on recent progress beating these bounds. We discuss how extremal bounds for set systems, when combined with other ideas, can improve on the Karger-Stein bound.

This is joint work with Euiwoong Lee (NYU) and Jason Li (CMU).

Thursday 31 October - Hervé Moulin (University of Glasgow)

Guarantees in Fair Division, under informational parsimony
Steinhaus's Diminishing Share (DS) algorithm (generalizing Divide & Choose D&C), as well as Dubins and Spanier's Moving Knife (MK) algorithm, guarantee to all participants a Fair Share of the manna (its worth at least 1/n-th of that of the whole

manna) while eliciting parsimonious information from them. However DS and MK are only defined when 1. preferences are represented by additive utilities; and 2. every part of the manna to be divided is desirable to every participant (a cake), or every part is unpleasant to everybody (a chore).

Our n-person Divide & Choose rule takes care of issue 2 when utilities are additive: it requires no trimming or padding, and works for mixed manna with subjective goods and bads. It also implements the canonical approximation of the Fair Share (up to one item) when we allocate indivisible items.

Issue 1 is much deeper, it challenges us to define a Fair Share Guarantee when 1/n-th of the whole manna makes no sense. The same D&C rule implements such a bound, for very general preferences restricted by a continuity assumption but no monotonicity whatsoever. The minMax utility of an agent is that of his best share in the worst possible partition. It is lower than his Maxmin utility (that of his worst share in the best possible partition), that cannot be guaranteed to all agents.

When the manna contains only goods, or only bads, the minMax Guarantee can be improved in infinitely many ways. Our Bid & Choose rules improve upon the MK rules by fixing a benchmark value of shares, and asking agents to bid the smallest size of an acceptable share. The resulting Guarantees fall between their minMax and Maxmin.

Joint work with Anna Bogomolnaia.

Friday 25 October - Liana Yepremyan (LSE)

On the size Ramsey number of graphs with bounded degree and bounded treewidth
A graph G is Ramsey for a graph H if every red/blue edge-colouring of the edges of G contains a monochromatic copy of H. We consider the following question: if H has bounded treewidth, is there a `sparse' graph G that is Ramsey for H? We show that if the maximum degree and treewidth of H are bounded, then there is a graph G with $O(|V(H)|)$ edges that is Ramsey for H. This was previously known for the smaller class of graphs H with bounded bandwidth by the work of by Clemens, Jenssen, Kohayakawa, Morrison, Mota, Reding and Roberts. We actually prove a more general off-diagonal version of the above result: For graphs H_1 and H_2, the \emph{size Ramsey number} $\hat{r}(H_1, H_2)$ is the minimum number of edges in a graph G such that every red/blue-colouring of the edges of G contains a red copy of H_1 or a blue copy of H_2. We prove that if H_1 and H_2 both have n vertices, bounded degree and bounded treewidth, then $\hat{r}(H_1, H_2)= O(n)$.

This is joint work with Nina Kam\v{c}ev, Anita Liebenau and David Wood.

Friday 18 October - Oliver Janzer (University of Cambridge)

The extremal number of subdivisions
For a graph H, the extremal number ex(n,H) is defined to be the maximal number of edges in an H-free graph on n vertices. For bipartite graphs H, determining the order of magnitude of ex(n,H) is notoriously difficult. In this talk I present recent progress on this problem.
The k-subdivision of a graph F is obtained by replacing the edges of F with internally vertex-disjoint paths of length k+1. Most of our results concern the extremal number of various subdivided graphs, especially the subdivisions of the complete graph and the complete bipartite graph.

Partially joint work with David Conlon and Joonkyung Lee.

Thursday 17 October - Anurag Bishnoi (FU Berlin)

Clique-free pseudorandom graphs
One of the outstanding open problems in the theory of pseudorandom graphs is to find a construction of K_s-free (n, d, λ)-graphs, for $s > 3$, with $\lambda = O(\sqrt{d})$ and the highest possible edge density of $d/n = \Theta(n^{-1/(2s - 3)})$.
For $s = 3$, there is a famous construction of Alon from 1994 that provides such a family of triangle free graphs.
For $s \geq 5$, the best known construction is due to Alon and Krivelevich from 1996 that has edge density $\Theta(n^{-1/(s - 2)})$.
Very recently, Mubayi and Verstraete have shown that a construction with edge density $\Omega(n^{-1/(s + \epsilon)})$, for any $\epsilon > 0$, would imply an improvement in the best known lower bounds on the off-diagonal Ramsey numbers $R(s, t)$, s fixed and $t \rightarrow \infty$.

In this talk I will present a new construction of K_s-free pseudorandom graphs with an edge density of $\Theta(n^{-1/(s - 1)})$, thus improving the Alon-Krivelevich construction but still falling short of improving the lower bounds on Ramsey numbers.

Joint work with Ferdinand Ihringer and Valentina Pepe.

Thursday 10 October - Ahmad Abdi (LSE)

Graphs, Matroids and Clutters (talk 2)
In a series of two talks, I will try to motivate and describe my area of research, and the mathematical objects that I deal with on a daily basis. The talks will be self-contained and will only assume basic knowledge of Linear Algebra, Polyhedral Theory, and Graph Theory.

The two talks center around the following conjecture that I made together with Gerard Cornuejols and Dabeen Lee:

"Let A be a 0,1 matrix where every row has at least two 1s and the polyhedron { x>=0 : Ax >= 1 } is integral. We conjecture that the columns of A can be partitioned into 4 color classes such that every row gets two 1s with different colors. This is still open even if 4 is replaced by any universal constant."

In the first talk, I will give two other equivalent formulations of this conjecture, one being the blocking version of this conjecture, the other being the "cuboidal" version.

In the second talk, I will talk about how this conjecture extends known prominent results in Graph Theory and Matroid Theory. In particular, we will see how the conjecture extends Jaeger's 8-flow theorem, and how a variation of it extends Tutte's 4-flow conjecture.

Friday 4 October - Natalie Behague (QMUL)

Semi-perfect 1-factorizations of the Hypercube
A 1-factorization of a graph H is a partition of the edges of H into disjoint perfect matchings {M_1 , M_2 , . . . , M_n}, also known as 1-factors. A 1-factorization M = {M_1 , M_2 , . . . , M_n} of a graph G is called perfect if the union of any pair of distinct 1-factors M_i , M_j is a Hamilton cycle. The existence or non-existence of perfect 1-factorizations has been studied for various families of graphs. Perhaps the most famous open problem in the area is Kotzig's conjecture, which states that the complete graph K_2n has a perfect 1-factorization. In this talk we shall focus on another well-studied family of graphs: the hypercubes Q_d in d dimensions. There is no perfect 1-factorization of Q_d for d > 2. As a result, we need to consider a weaker concept.
A 1-factorization M is called k-semi-perfect if the union of any pair of 1-factors M_i , M_j with 1 ≤ i ≤ k and k + 1 ≤ j ≤ n is a Hamilton cycle. It was proved that there is a 1-semi-perfect 1-factorization of Q_d for every integer d ≥ 2 by Gochev and Gotchev, Královič and Královič, and Chitra and Muthusamy, in answer to a conjecture of Craft. My main result is a proof that there is a k-semi-perfect 1-factorization of Q_d for all k and all d, except for one possible exception when k = 3 and d = 6. I will sketch the proof and explain why this is, in some sense, best possible. I will conclude with some questions concerning other generalisations of perfect 1-factorizations.

Thursday 3 October - Ahmad Abdi (LSE)

Clutters, blockers, and cuboids (talk 1)
In a series of two talks, I will try to motivate and describe my area of research, and the mathematical objects that I deal with on a daily basis. The talks will be self-

contained and will only assume basic knowledge of Linear Algebra, Polyhedral Theory, and Graph Theory.

The two talks center around the following conjecture that I made together with Gerard Cornuejols and Dabeen Lee:

"Let A be a 0,1 matrix where every row has at least two 1s and the polyhedron { x>=0 : Ax >= 1 } is integral. We conjecture that the columns of A can be partitioned into 4 color classes such that every row gets two 1s with different colors. This is still open even if 4 is replaced by any universal constant."

In the first talk, I will give two other equivalent formulations of this conjecture, one being the blocking version of this conjecture, the other being the "cuboidal" version.

In the second talk, I will talk about how this conjecture extends known prominent results in Graph Theory and Matroid Theory. In particular, we will see how the conjecture extends Jaeger's 8-flow theorem, and how a variation of it extends Tutte's 4-flow conjecture.

Wednesday 12 June - Benny Sudakov (ETH)

Covering graphs by monochromatic trees and Helly-type results for hypergraphs
How many monochromatic paths, cycles or general trees does one need to cover all vertices of a given r-edge-coloured graph G? These problems were introduced in the 1960s and were intensively studied by various researchers over the last 50 years. In this talk, we establish a connection between this problem and the following natural Helly-type question in hypergraphs. Roughly speaking, this question asks for the maximum number of vertices needed to cover all the edges of a hypergraph H if it is known that any collection of a few edges of H has a small cover. We obtain quite accurate bounds for the hypergraph problem and use them to give some unexpected answers to several questions about covering graphs by monochromatic trees raised and studied by Bal and DeBiasio, Kohayakawa, Mota and Schacht, Lang and Lo, and Girao, Letzter and Sahasrabudhe.

Joint work with M. Bucic and D. Korandi

Wednesday 22 May - R. Ravi (CMU)

Hub Network Design
We study a fundamental class of hub network design problems that arise in the design of logistic networks that are configured in two layers. An upper layer is configured by establishing hubs at selected nodes at considerable cost so that the routes between hubs can be operated cheaply. The remaining edges in the network are operated at regular cost. The resulting problem is to determine the set of nodes

to open hubs so that the edges induced between them have cheaper costs and the goal is to find a network of minimum total edge plus hub opening costs.
We study several hub network design problems such as spanning trees, Steiner trees, generalized Steiner forests and further one-connected generalizations modelled as proper function cut covers. When edge costs are non-metric, we show logarithmic approximation hardness even for spanning trees; On the other hand, we design approximation algorithms with matching logarithmic performance ratios for the most general one-connected proper function cut-cover problems in this setting including Steiner trees and generalized Steiner forests. These results generalize and use the idea of spider decompositions introduced for node-weighted network design problems. When the costs are metric, we develop constant-factor approximation algorithms using a truncated primal-dual approach for the hub Steiner tree problem.

This is joint work with Takuro Fukunaga (Chuo University, Japan), Oleksandr Rudenko and Ziye Tang(CMU)

Wednesday 15 May - Sergey Sergeev (Birmingham)

Tropical Optimisation and Mean-Payoff Games
The connection between tropical polyhedra and mean-payoff games, established in a work of Akian, Gaubert and Guterman, can be also used to solve a number of optimisation problems over tropical polyhedra. We are going to consider, in particular, the problems of tropical linear-fractional and tropical pseudo-quadratic optimisation and discuss how they can be solved by Newton and bisection schemes involving mean-payoff game solvers.

Monday 29 April - Christoph Hunkenschröder (EPFL)

Integer Programming from the viewpoint of Parameterized Complexity
In the 1980’s, Lenstra showed that Integer Programming can be solved in polynomial time if the dimension is fixed. In this talk, we will consider the dimension to be part of the input, and focus on certain other parameters.
In particular, given an Integer Program in equality form Ax = b with integral matrix A, we will introduce the tree-depth of the (primal or dual) constraint graph of A as a measure of density for A. If the entries of A and the tree-depth are small, we can solve Integer Programming efficiently, in particular in time polynomial in the dimension.

This is joint work with Friedrich Eisenbrand, Kim-Manuel Klein, Martin Koutecký, Asaf Levin and Shmuel Onn.

Thursday 2 May - Dona Strauss (University of Leeds)

Subsets of βN which are not Borel
Many subsets of βN which are simple to define algebraically, are far from simple topologically. I shall show that the following subsets of βN are not Borel: the set of idempotents; the smallest ideal; any principal right ideal defined by an element of the remainder space N^*; N^*+N^*.

Thursday 28 March - Marcin Jurdzinski (University of Warwick)

Universal trees and quasi-polynomial bounds for parity games
Several distinct techniques have been proposed to solve parity games in quasi-polynomial time since the breakthrough result of Calude, Jain, Khoussainov, Li, and Stephan (2017): play summaries, progress measures, and register games. We argue that all those techniques can be viewed as instances of the separation approach to solving parity games---a key component of which is constructing an automaton that separates languages of words encoding plays that are decisively won by either of the two players---thus establishing a single model that unifies the recent approaches to solving parity games efficiently.
Our main technical results are the nearly matching quasi-polynomial upper and lower bounds on the sizes of all separating automata. In particular, the lower bound forms a barrier that the existing approaches must overcome in the ongoing quest for a polynomial-time algorithm for solving parity games.We obtain our main results by introducing and studying universal ordered trees, a combinatorial object that may be of independent interest. Firstly, we establish the nearly matching quasi-polynomial upper and lower bounds on the size of smallest universal trees. Then we prove that the sizes of the smallest separating automata and of the smallest universal trees coincide---and hence both are quasi-polynomial---by showing how to turn universal as trees into separating automata, and that there is a universal tree hiding in the states of every separating automaton.

The talk is based on joint work with Wojciech Czerwiński, Laure Daviaud, Nathanaël Fijalkow, Ranko Lazić, and Paweł Parys.

Thursday 21 March - Adrian Mathias

Flutters and chameleons
Let K be the collection of infinite sets of natural numbers. A colouring c of K with a finite number of colours is Ramsey if for some infinite A in K and every infinite subset B of A, c(B) = c(A). A non-Ramsey colouring is one for which no such A exists.

Solovay in a famous paper published in 1970 used a strongly inaccessible cardinal to construct a model of ZF + DC in which these principles hold which contradict AC:

LM: every set of real numbers is Lebesgue measurable;
PB: every set of real numbers has the property of Baire;
UP: every uncountable set of real numbers has a perfect subset.

Two other principles to be considered are
RAM: all colourings are Ramsey

NoMAD: there is no maximal infinite family of pairwise almost disjoint infinite sets of natural numbers.
The speaker showed in 1968 that in Solovay's model, RAM holds, and in 1969 that if one started from a Mahlo cardinal, NoMAD would hold in the corresponding Solovay model. It is natural to ask whether these large cardinals are necessary; the inaccessible is necessary for UP (Specker) and LM (Shelah) but not for PB (Shelah).

More recently Toernquist has shown that NoMAD holds in Solovay's original model, and Shelah and Horowitz have extended his work to show that even that inaccessible is unnecessary to get a model of NoMAD. But it has been open for fifty years whether RAM requires an inaccessible. This talk will be chiefly about flutters and chameleons, which are non-Ramsey sets with elegant properties, constructed using weak forms of AC; surprisingly their existence has been found to follow from various Pareto principles of mathematical economics.

Wednesday 20 March - Davi Silva (Cologne)

Stronger transference principles: preserving more structure
Transference principles (also called dense model theorems) are results which allow us to transfer some combinatorial theorems from the usual positive-density setting to the very sparse setting, provided the sparse objects satisfy some mild boundedness conditions.
In this talk I will give an overview of these results and the main ideas and concepts behind them. I will then explain in which way they can be made stronger, and give some preliminary results obtained in this direction.

Thursday 14 March - Jugal Garg (ISE)

A Strongly Polynomial Algorithm for Linear Exchange Markets
In this talk, I will present the first strongly polynomial algorithm for computing an equilibrium in exchange markets with linear utilities. The exchange market model has been extensively studied since its introduction by Walras in 1874. This problem has a non-separable convex flow formulation and the property that we can find an equilibrium in strongly polynomial time given its support, i.e., the flow variables which are non-zero. Using a variant of Duan and Mehlhorn (DM) algorithm, we gradually reveal new variables that are in the support of every equilibrium. We show that a new variable can be revealed in strongly polynomial time if we start the DM

algorithm with the best possible solution corresponding to the current revealed set. Finding the best solution can be reduced to solving a linear program (LP). Even though we are unable to solve this LP in strongly polynomial time, we show that it can be approximated by a simpler LP with two variables per inequality that is solvable in strongly polynomial time and it turns out to be good enough to make the overall algorithm run in strongly polynomial time. This is based on joint work with Laci Vegh.

Thursday 7 March - Simone Cerreia Vioglio (Bocconi)

Robust Opinion Aggregation and its Dynamics
We consider a robustification of the DeGroot's linear model of social learning. By relaxing the assumption of quadratic utility of the agents, we obtain an opinion aggregator that is normalized, monotone, and translation invariant. We directly link these properties to natural conditions of the microfoundation. In addition to the less demanding assumptions on the payoff function of the agents, the opinion aggregator allows for several economically relevant patterns ruled out by the more restrictive linear model. For instance, agents can feature homophily, dislike (or attraction) for extreme opinions as well as discard information obtained from sources that are perceived as redundant. We also show that under these weaker assumptions is still possible to explore the standard questions addressed by the linear model, such as convergence of limit opinions, and the properties of consensus and wisdom for this limit.

Joint work with R. Corrao (MIT) and G. Lanzani (MIT).

Wednesday 6 March - Dani Dorfman (TAU)

A faster deterministic exponential time algorithm for Energy Games and Mean Payoff Games
We present an improved exponential time algorithm for Energy Games, and hence also for Mean Payoff Games. The running time of the new algorithm is O(min(m*n*W, n*m* 2^(n/2)* logW), where n is the number of vertices, m is the number of edges, and when the edge weights are integers of absolute value at most W. The new algorithm is also a pseudopolynomial time algorithm, matching the O(mnW) running time of the fastest known pseudopolynomial time algorithm of Brim et al. on which it is based.It is currently the fastest deterministic algorithm for Energy Games and Mean Payoff Games when 2^(n/2})*n < W < 2^(2^(n/2)). The new algorithm is obtained by introducing a technique of forecasting repetitive actions performed by the algorithm of Brim et al., along with the use of an edge-weight scaling technique.

Thursday 28 February - Yani Pehova (Warwick)

Tilings in graphs with sublinear independence number
A classical theorem of Hajnal and Szemeredi states that if an n-vertex graph G has minimum degree at least (1-1/r)n, then it contains a K_r-factor (that is, n/r vertex-disjoint copies of K_r), provided r divides |G|. Extremal examples for this result, however, contain large independent sets. Forbidding these extremal structures is likely to decrease the minimum degree condition required to force a K_r-factor, and indeed a result of Balogh, Molla and Sharifzadeh shows that if the independence number of G is sublinear, then minimum degree slightly above n/2 suffices to force a triangle-factor. We present an extension of this result for general K_r-factors and graphs of sublinear k-independence number (that is, whose largest K_k-free subsets have sublinear size). Our proof uses the absorbing method, and our absorber construction can be applied in other settings. For example, we are able to recover a recent result of Balogh, Treglown and Wagner on H-factors in dense randomly perturbed graphs. This is joint work with Rajko Nenadov.

Wednesday 27 February - Georg Loho (LSE)

Algorithmic questions around tropical Carathéodory
Since Imre Bárány found the colourful version of Carathéodory's theorem in 1982, many combinatorial generalizations and algorithmic variations have been considered. This ranges from variations of the colour classes to different notions of convexity. We take a closer look at the max-plus or 'tropical' convexity version of this theorem. We provide new insights on colourful linear programming and matroid generalizations from a tropical point of view, by considering additional sign informations. The difficulty of the arising algorithmic questions ranges from greedily solvable to NP-hard.

Thursday 14 February - Cécile Mailler (Bath)

The monkey walk: a random walk with random reinforced relocations and fading memory.
In this joint work with Gerónimo Uribe-Bravo, we prove and extend results from the physics literature about a random walk with random reinforced relocations. The "walker" evolves in $\mathbb Z^d$ or $\mathbb R^d$ according to a Markov process, except at some random jump-times, where it chooses a time uniformly at random in its past, and instatnly jumps to the position it was at that random time. This walk is by definition non-Markovian, since the walker needs to remember all its past.
Under moment conditions on the inter-jump-times, and provided that the underlying Markov process verifies a distributional limit theorem, we show a distributional limit theorem for the position of the walker at large time. The proof relies on exploiting the branching structure of this random walk with random relocations; we are able

to extend the model further by allowing the memory of the walker to decay with time.

Wednesday 13 February - Stefano Leonardi (Sapienza Università di Roma)

Algorithmic Mechanism Design for Two-sided Markets
Mechanism design for one-sided markets is an area of extensive research in economics and, since more than a decade, in computer science as well. Two-sided markets, on the other hand, have not received the same attention despite the many applications to Internet advertisement and to the sharing economy.
In two-sided markets, both buyers and sellers act strategically. An ideal goal in two-sided markets is to maximize the social welfare of buyers and sellers with individually rational (IR), incentive compatible (IC) and budget-balanced mechanisms (BB), which requires that the mechanism does not subsidize the market or make an arbitrary profit from the exchange. Unfortunately, Myerson and Satterthwaite proved in 1983 that this goal cannot be achieved even in the bayesian setting and for the simple case of only one buyer and one seller.
In this talk, I'll discuss meaningful trade-offs and algorithmic approximations of the above requirements by presenting several recent results and some challenging open problems.

Thursday 7 February - Carla Groenland (University of Oxford)

Size reconstructability of graphs
The deck of a graph G is given by the multiset $\{G-v:v\in V(G)\}$ of (unlabelled) subgraphs, which are called cards. The graph reconstruction conjecture states that no two non-isomorphic graphs (on at least three vertices) can have the same deck. The number of edges of a graph can be computed from its deck and Brown and Fenner show that the size of G can be reconstructed as well after any 2 cards have been removed from the deck (for $n\geq 29$). We show that for sufficiently large n, the number of edges can be computed whenever at most $\frac1{20}\sqrt{n}$ cards are missing. I will talk about the background of the problem, the main ideas of the proof and my favourite open problems in the area.Joint work with Hannah Guggiari and Alex Scott.

Wednesday 30 January - Victor Verdugo (LSE)

Breaking symmetries to rescue Sum of Squares
The Sum of Squares (SoS) hierarchy gives an automatized technique to create a family of increasingly tight convex relaxations for binary programs. There are several problems for which a constant number of rounds of the hierarchy give integrality gaps matching the best known approximation algorithm. In many other, however, ad-hoc techniques give significantly better approximation ratios. Notably, the lower bounds instances, in many cases, are invariant under the action of a large

permutation group. In the first part of the talk I'll provide some background about polynomial optimization and the SoS approach. In the second part, I'll show how symmetries can be broken in order to obtain arbitrarily good approximation ratios for the case of makespan scheduling. On the negative side, SoS fails when symmetries are not broken, and I'll highlight the ideas behind the lower bound.

Wednesday 23 January - Peter Allen (LSE)

Optimal packings of sparse graphs
I will describe (another) randomised approach to finding perfect packings of sparse graphs in quasirandom or complete host graphs. The packing algorithm is simple and I will describe it completely; the analysis is not so simple. In particular, we prove that Ringel's Conjecture and the Tree Packing Conjecture are true for almost all trees (or sequences of trees, respectively).

This is joint work with Julia Boettcher, Dennis Clemens, Jan Hladky, Diana Piguet and Anusch Taraz.