Operations Research Seminar Series in 2014

Seminars are listed in reverse chronological order, most recent first.

---

10 December - Alec Waterhouse (Department of Energy and Climate Change)
Modelling in Department of Energy Climate Change an Overview
Alec Waterhouse the Head of the Central Modelling Team in DECC will give a short introduction to the analytical modelling work of the Department of Energy and Climate Change. He will describe DECC's objectives and give an overview of modelling landscape and some examples of the types of models used in the Department. He will put these models into the context of the overall modelling strategy, stress the importance of quality assurance for anyone wishing to interact with government departments and outline some of the challenges faced by the department.

3 December - Daniel Dadush (CWI Amsterdam)
New Bounds for Curved Polyhedra via the Shadow Simplex Method

We study the simplex method over polyhedra satisfying certain "discrete curvature" lower bounds. Intuitively speaking, for the polyhedra in question, we enforce that the boundary always meets vertices at ``sharp'' angles.
Our work builds and improves upon the results of Bonifas et al (SOCG 2012), Brunsch and Roglin (ICALP 2013), and Eisenbrand, Vempala (2014). As our main results, we give an improved (constructive) diameter bound over these polytopes, as well as a faster simplex based algorithm for linear optimization. As our main technical tool, we develop a new analysis and variant of the shadow simplex method.

More precisely, for an n-dimensional polyhedron with m facets and curvature parameter 0 < delta < 1, we give a diameter bound of O(n^2(1+ ln(n/delta))/delta). For the class of polyhedra having totally unimodular constraint matrices, this implies an O(n^3 ln n) diameter bound. For linear optimization, given an initial feasible vertex, we show that an optimal vertex can be found using an expected O(n^3(1+ln(n/delta))/delta) simplex pivots, each requiring O(m n) time to compute. Furthermore, a first feasible solution can be found using O(m n^3(1+ln(n/delta))/delta) pivot steps.

This is joint work with Nicolai Hahnle (University of Bonn).

26 November - Gilad Bavly (Tel Aviv)
Elasticity of games

We develop an elasticity index of a strategic game. The index measures the robustness of the set of rational outcomes of a game. The elasticity index of a game is the maximal ratio between the change of the rational outcomes and the size of an infinitesimal perturbation. The perturbation is on the players' knowledge of the game.

The elasticity of a strategic game is a nonnegative number. A small elasticity is indicative of the robustness of the rational outcomes (for example, if there is only one player the elasticity is 0), and a large elasticity is indicative of non-robustness. For example, the elasticity of the (normalized) n-stage finitely repeated prisoner's dilemma is at least exponential in n, as is the elasticity of the n-stage centipede game and the n-ranged traveler's dilemma.

The concept of elasticity enables us to look from a different perspective at Neyman's (1999) repeated games when the number of repetitions is not commonly known, and Aumann's (1992) demonstration of the effect of irrationality perturbations.

19 November - Imre Bárány (UCL/Rényi Institute)
Block partitions of sequences

Given a sequence A=(a_1,...,a_n) of real numbers, a block B of A is either a set B={a_i,a_{i+1},...,a_j} where i <= j or the empty set. The size b of a block B is the sum of its elements. We show that when each a_i lies in [0,1] and k is a positive integer, then there is a partition of A into k blocks B_1,...,B_k such that |b_i-b_j| is at most one for every i,j. We extend this result in several directions.
This is joint work with Victor Grinberg.

5 November - John Howard (LSE)
Strategic Optimisation of Rank

In many competitive situations (including nearly all sports) a player's aim is not simply to maximize his score but to maximize its rank among all scores. Examples include sales contests (where the salesman with highest monthly sales gets a bonus) and patent races (where lowest time is best). We assume a player's score is obtained costlessly, so that his utility is the probability of having the best score. This gives a constant-sum game. All that matters for a player is the distribution of his score, so we assume he chooses from a given convex set of distributions. We prove that such games, called Distribution Ranking Games, under certain assumptions have an equilibrium solution in pure strategies. We characterize their

solution for various classes of distributions, such as distributions with given mean or moment, where we extend a result of Bell and Cover. Our model was stimulated by an intriguing game of E. J. Anderson: we modify his game to a setting in which several players go into separate casinos with given starting capital, and the one who stops with the highest amount wins.

This is joint work with Professor Steve Alpern.

22 October - Alina R. Ene (Warwick)
Approximation algorithms for multiway partitioning problems and a simplex coloring conjecture

We consider several problems where the goal is partition a ground set into several pieces while minimizing a "cut-type" objective function; examples include Multiway Cut, Node-weighted Multiway Cut, Metric Labeling and Hypergraph Labeling. A natural LP relaxation gives an optimal approximation for these problems, assuming the Unique Games Conjecture (the UGC assumption can be removed for certain submodular generalizations of these problems). However, we do not know how to round this LP in general and the focus has been on understanding this LP for specific problems. In this talk, we describe several rounding strategies and an integrality gap construction that leads to a simplex coloring conjecture reminiscent of Sperner's Lemma.

This talk is based on joint work with Chandra Chekuri (UIUC), Huy Nguyen (Simons Institute Berkeley), Jan Vondrak (IBM Research Almaden), and Yi Wu (Google).

15 October - Arndt von Schemde (Thema Consulting Group)
Power market modelling using linear programming

Power markets can be simulated using a linear programming approach. The intersection between the demand and the supply curves is found by solving a cost minimization problem under a set of constraints. In this talk, we give an example of a power market model (the "TheMA" model) and its applications. The model has an hourly time resolution, a detailed representation of thermal, hydro, and intermitted generation technologies (such as wind and PV), as well as a detailed description of cross-border trade. In this talk, we shall also discuss practical considerations, such as computation time issues, and explain what general principles should guide modelling and model development.

8 October - Katerina Papadaki (LSE)
Patrolling games

A key operational problem for those charged with the security of vulnerable facilities (such as airports or art galleries) is the scheduling and deployment of

patrols. Motivated by the problem of optimizing randomized, and thus unpredictable, patrols, we present a class of patrolling games. The facility to be patrolled can be thought of as a network or graph Q of interconnected nodes (e.g., rooms, terminals), and the Attacker can choose to attack any node of Q within a given time T. He requires m consecutive periods there, uninterrupted by the Patroller, to commit his nefarious act (and win). The Patroller can follow any path on the graph. Thus, the patrolling game is a win-lose game, where the Value is the probability that the Patroller successfully intercepts an attack, given best play on both sides. We determine analytically either the Value of the game, or bounds on the Value, for various classes of graphs, and we discuss possible extensions and generalizations.

2 July - Evdokia Nikolova (Austin)
Approximation Algorithms for Offline Risk-averse Combinatorial Optimization
We consider generic optimization problems that can be formulated as minimizing the cost of a feasible solution w.x over a combinatorial feasible set F ⊂ {0, 1}^n. For these problems we describe a framework of risk-averse stochastic problems where the cost vector W has independent random components, unknown at the time of solution. A natural and important objective that incorporates risk in this stochastic setting is to look for a feasible solution whose stochastic cost has a small tail or a small convex combination of mean and standard deviation. Our models can be equivalently reformulated as nonconvex programs for which no efficient algorithms are known.
We provide efficient general-purpose approximation algorithms. They use as a black-box (exact or approximate) the solution to the underlying deterministic problem and thus immediately apply to arbitrary combinatorial problems. For example, from an available approximation algorithm to the deterministic problem, we construct an approximation algorithm with almost the same approximation factor for the stochastic problem. The algorithms are based on a geometric analysis of the nonlinear level sets of the objective functions.

18 June - Shmuel Gal (Haifa)
Succession of hide-seek and pursuit-evasion at heterogeneous locations

A predator (searcher) looks for a prey (hider) in a search space consisting of n locations. The hider chooses a location and the searcher inspects k different locations, where k is a parameter of the game (the `giving-up time' for the continuous version). If the predator visits a location i at which the prey hides, then the game moves into a pursuit-evasion phase. In this phase capture is not certain but occurs with probability p_i.

We show that for all k smaller than an easily calculated threshold, it is optimal to hide with probability proportional to 1/p_i for each location i: If k exceeds the threshold, then the optimal hiding strategy is always to stay at the location with the

smallest p_i.

We extend this game to a repeated game. During the k looks among the different locations within a single patch, there can be any of three events. First, if the searcher does not find the hider, then the game ends with zero payoff for the searcher and a payoff of one to the hider. Second, if the searcher finds the hider and catches it, then the game ends with a payoff of one to the searcher and zero to the hider. Finally, if the searcher finds the hider but does not catch it then the hider escapes to another patch and the process restarts. We show that in this game the optimal hiding strategy is to always make all the locations equally "attractive" for the searcher, no matter how large is k: This situation is quite different from the one stage game in which solutions of this type occur only if k is below the threshold.

11 June - Christopher S. Tang (Edward W. Carter Professor of Business Administration, UCLA Anderson School)
Project Design with Limited Commitment and Teams

We study the interaction between a group of agents who exert costly effort over time to complete a project, and a manager who chooses the objectives that must be met in order for her to sign off on it. The manager has limited commitment power so that she can commit to the requirements only when the project is sufficiently close to completion. This is common in projects that involve design or quality objectives, which are difficult to define far in advance. The main result is that the manager has incentives to extend the project as it progresses: she is time-inconsistent. This result has two implications. First, the manager will choose a larger project if she has less commitment power. Second, if the agents receive a fraction of the project's worth upon its completion, then the manager should delegate the decision rights over the project size to the agents unless she has sufficient commitment power. In this case, the agents will choose a smaller project than is optimal for the manager, but their preferences are time-consistent.
This is joint work with George Georgiadis and Steven Lippman.

4 June - Spyros Angelopoulosk (University Pierre et Marie Curie)
Star-search problems: new measures, algorithms and techniques

We consider the problem of exploring a set of m concurrent rays using a single searcher. The rays are disjoint with the exception of a single common point, and in each ray a potential target may be located. The objective is to design efficient search strategies for locating the targets as quickly as possible.
In this talk I will describe results on two variants of this well-studied problem. In the first variant, the searcher must locate t targets (with t less than or equal to m); this generalizes the setting in which only a single target is sought. The second variant corresponds to the setting in which the searcher incurs a fixed cost upon switching

direction. For this problem, I will describe a proper use of infinite LP formulations that yields tight lower bounds.

27 May - Robbert Fokkink (TU Delft)
The search value of a set

In 1985, Alpern and Asic defined the search value of a network as the expected time of finding a hider in a search game, when time is scaled such that an exhaustive search takes unit time. The search value of a network is always in between ½ and 1. It has not been computed for many networks yet. I will try to explain why this is so. Extending the Alpern-Asic concept, one can define the search value of a set. This is remotely related to the Shapley value of a cooperative game. Again, the search value of a set is in between ½ and 1. Again, it is hard to compute. I will give you one such set, in which searching comes at a supermarket discount of 5 for the price of 3, and leave the computation of its value as a challenge.

This is joint work with Ken Kikuta (Kobe) and David Ramsey (Wroclaw).

21 May - Adam Letchford (Lancaster)
Stronger multi-commodity flow formulations for the capacitated vehicle routing problem

The capacitated vehicle routing problem (CVRP) is a much-studied NP-hard combinatorial optimization problem. Many different integer programming formulations have been proposed for the CVRP, including so-called single-commodity flow (SCF) and multi-commodity flow (MCF) formulations. We review these formulations, and then present two new MCF formulations. We show that they dominate all of the existing SCF and MCF formulations, in the sense that their continuous relaxations yield stronger lower bounds. Some preliminary computational results will also be presented. This talk is based on joint work with Professor Juan-Jose' Salazar from the University of La Laguna.

30 April - Stanislav Živný (Oxford)
The complexity of finite-valued CSPs

Let L be a set of rational-valued functions on a fixed finite domain; such a set is called a finite-valued constraint language. We are interested in the problem of minimising a function given explicitly as a sum of functions from L. We establish a dichotomy theorem with respect to exact solvability for all finite-valued languages defined on domains of arbitrary finite size. We present a simple algebraic condition that characterises the tractable cases. Moreover, we show that a single algorithm based on linear programming solves all tractable cases. Furthermore, we show that there is a single reason for intractability; namely, a very specific reduction from Max-Cut.

This is joint work with J. Thapper.

19 March - Matthias Mnich (TU Darmstadt)
Improved Integrality Gap Upper Bounds for TSP with Distances One and Two

We study the structure of solutions to linear programming formulations for the traveling salesperson problem (TSP). We perform a detailed analysis of the support of the subtour elimination linear programming relaxation, which leads to algorithms that find 2-matchings with few components in polynomial time. The number of components directly leads to integrality gap upper bounds for the TSP with distances one and two, for both undirected and directed graphs.

Our main results for fractionally Hamiltonian instances are:

- For undirected instances we obtain an integrality gap upper bound of 5/4 without any restrictions, of 7/6 if the optimal LP solution is half-integral, and of 10/9 if there is an optimal solution that is a basic solution of the fractional 2-matching polytope.
- For directed instances we obtain an integrality gap upper bound of 3/2, and of 4/3 if given an optimal 1/2-integral solution. Our algorithms perform sequences of local improvements that harness the structure of the support.

Additionally, we show that relying on the structure of the support is not an artifact of our algorithm, but is necessary under standard complexity-theoretic assumptions: we show that finding improved solutions via local search is W[1]-hard for k-edge change neighborhoods even for the TSP with distances one and two, which strengthens a result of Dániel Marx.

This is joint work with Tobias Moemke.

12 March - Steven Alpern (Warwick Business School)
Expanding Search on Networks

This joint work with Tom Lidbetter considers searching a network in a new way, which we call "expanding search". In the discrete case, where the Hider is restricted to hiding at one of the k nodes of a network Q, such a search is a sequence of distinct nodes n(1), n(2),...,n(k), where n(1) is a given 'root' or 'starting' node, and every other node n(i) is adjacent to a previously searched node n(j), for some j<i. We consider both the Bayesian problem where the Hider distribution is known, and the search game where the Hider picks his node. The payoff is the expected search time. We also consider the case where the Hider can hide at any point on the network (where arcs have given lengths).

5 March - Frans de Ruiter (LSE)
Adjustable robust optimization with decision rules based on inexact revealed data

Adjustable robust optimization (ARO) is a technique to solve dynamic (multistage) optimization problems. In ARO, the decision in each stage is a function of the information accumulated from the previous periods on the values of the uncertain parameters. This information, however, is often inaccurate; there is much evidence in the information management literature that even in our Big Data era the data quality is often poor. Reliance on the data "as is" may then lead to poor performance of ARO, or in fact to any "data-driven" method. In this paper, we remedy this weakness of ARO by introducing a methodology that treats past data itself as an uncertain parameter. We show that algorithmic tractability of the robust counterparts associated with this extension of ARO is still maintained. The benefit of the new approach is demonstrated by a production-inventory application.

This is joint work with Aharon Ben-Tal (Technion, Haifa Israel), Dick den Hertog and Ruud Brekelmans (Tilburg University).

26 February - Greg Sorkin (LSE)
VCG Auction Mechanism Cost Expectations and Variances

We consider Vickrey--Clarke--Groves (VCG) auctions for a very general combinatorial structure, in an average-case setting where item costs are independent, identically distributed uniform random variables. We prove that the expected VCG cost is at least double the expected nominal cost, and exactly double when the desired structure is a basis of a bridgeless matroid. In the matroid case we further show that, conditioned upon the VCG cost, the expectation of the nominal cost is exactly half the VCG cost, and we show several results on variances and covariances among the nominal cost, the VCG cost, and related quantities. We provide examples showing that our results cannot be strengthened in some natural ways.

Our methods allow calculation of the VCG cost variance in some cases, including (asymptotically) for a minimum spanning tree in a complete graph with random edge costs.

This is joint work with Svante Janson (Uppsala).

12 February - Om Narasimhan (LSE)
Is Cash King? A Field Intervention on Mental Accounting in a Salesforce

The fungibility of money, along with its function as a medium of exchange, gives monetary transactions pride of place in neoclassical theories of economics.

However, a growing number of laboratory studies and anecdotal field evidence (appealing to theories of mental accounting) have shown that consumers' willingness to spend an additional unit of wealth depends upon the sources and categories of wealth, thus throwing serious doubt on the presumed fungibility of cash. This paper extends the work to examining the fungibility of money in the context of salesperson effort. Given that firms desire to design the least cost pay basket to evoke a given level of effort, it is important to i) identify whether salespersons differ in their response to disparate forms of compensation, ii) quantify the impact of different sources of wealth, and iii) explore the underlying mechanisms at play.

We seek to achieve the above objectives using a multi-method empirical approach. In particular, we use a field intervention, reduced form estimation to derive initial insights from the intervention, a structural model to uncover the differential response to cash versus non-cash compensation, and a survey of salespeople to pinpoint the precursors of separate budgets. We find converging evidence that salespeople maintain separate accounts for cash and non-cash bonuses, dispelling the neo-classical expectation that favors cash as king.

5 February - Rico Zenklusen (ETH Zürich)
Multi-Budgeted Matchings via the Ham Sandwich Theorem

In many applications, one has to deal with multiple, partially conflicting constraints. In this talk, we study a multi-objective variant of a classical combinatorial optimization problem, namely the maximum weight matching problem. A natural way to deal with several objectives is to turn all of the objectives but one into budget constraints. This leads to the multi-budgeted matching problem which asks to find a maximum weight matching subject to k linear constraints with nonnegative coefficients. Whereas this problem can easily be shown to be NP hard even for k=1, I will present in this talk a polynomial time approximation scheme that works for any constant k.

To prove that our algorithm is correct, we leverage two beautiful non-constructive mathematical theorems. More precisely, the Jordan Curve Theorem gives a concise and intuitive proof why our algorithm works for k=1, and a result of Stromquist and Woodall that follows from the Ham Sandwich Theorem allows for showing correctness for any constant k.

Part of this work is joint with Fabrizio Grandoni.

29 January - Bernhard von Stengel (LSE)
Recursive Inspection Games

Dresher (1962) described a sequential inspection game where an inspector has to distribute a given number of inspections over a larger number of inspection periods in order to detect an illegal act that an inspectee, who can count the inspector's visits, performs in at most one of these periods. We present an extension of this game where more than one illegal act is allowed. Then, under certain reasonable assumptions for the zero-sum payoffs, the optimal strategy of the inspector does not depend on the number of intended illegal acts. This allows a recursive description. The resulting recursive equation in three variables for the value of the game, which generalizes several other known equations of this kind, is solved explicitly, using numbers similar to binomial coefficients defined by a "generalized Pascal triangle". We also extend this approach to non-zero-sum games and, similar to Maschler (1966), "inspector leadership" where the inspector commits to (the same) randomized inspection schedule, but the inspectee acts legally as long as inspections remain.

22 January - Magnus Wahlström (Royal Holloway)
Half-integrality, LP-branching and FPT Algorithms

A recent trend in parameterized algorithms is the application of polytope tools (specifically, LP-branching) to FPT algorithms (e.g., Cygan et al., 2011; Narayanaswamy et al., 2012). Though the list of work in this direction is short, the results are already interesting, yielding significant speedups for a range of important problems. However, the existing approaches require the underlying polytope to have very restrictive properties, including half-integrality and Nemhauser-Trotter-style persistence properties. To date, these properties are essentially known to hold only for two classes of polytopes, covering the cases of Vertex Cover (Nemhauser and Trotter, 1975) and Node Multiway Cut (Garg et al., 1994).

Taking a slightly different approach, we view half-integrality as a discrete relaxation of a problem, e.g., a relaxation of the search space from {0,1\}^V to {0,1/2,1}^V such that the new problem admits a polynomial-time exact solution. Using tools from CSP (in particular Thapper and Zivny, 2012) to study the existence of such relaxations, we are able to provide a much broader class of half-integral polytopes with the required properties.

Our results unify and significantly extend the previously known cases. In addition to the new insight into problems with half-integral relaxations, our results yield a range of new and improved FPT algorithms, including an O*(|Sigma|^{2k})-time algorithm for node-deletion Unique Label Cover with label set Sigma (improving the previous bound of O*(|\Sigma|^{O(k^2 \log k)}) due to Chitnis et al. (2012) and an O*(4^k)-

time algorithm for Group Feedback Vertex Set, including the setting where the group is only given by oracle access (improving on the previous bound of $O^*(2^{O(k\log k)})$ due to Cygan et al. (2012). The latter bound is optimal under the Exponential Time Hypothesis. The latter result also implies the first single-exponential time FPT algorithm for Subset Feedback Vertex Set, answering an open question of Cygan et al. (2012).

Interestingly, despite the half-integrality, our result do not imply any approximation results (as may be expected, given the Unique Games-hardness of the covered problems).