Thursday, 15 July 2010

Talks by Participants		
9:15 – 9:45	Sascha Geulen (RWTH Aachen)	Online Capacity Maximization in Wireless Networks
9:45 – 10:15	Harald Räcke (University of Warwick)	Oblivious Routing in the L_p-norm
10:15 – 10:45	Rahul Savani (University of Liverpool)	The Complexity of the Homotopy Method, Equilibrium Selection, and Lemke-Howson Solutions
10:45 – 11:15		Coffee Break
11:15 – 11:45	Tobias Brunsch (Maastricht University)	Clustering - How Bad Is The k-Means++ Method?
11:45 – 12:15	Martin Gairing (University of Liverpool)	Covering Games: Approximation Through Non-Cooperation
12:15 – 12:45	Rajiv Raman (University of Warwick)	Maximum Feasible Subsystems
12:45 – 14:45		Lunch Break
Afternoon Session		

Friday, 16 July 2010

Talks by Participants		
9:15 – 9:45	Heiko Röglin (Maastricht University)	Smoothed Analysis of Multiobjective Optimization
9:45 – 10:15	Mike Paterson (University of Warwick)	Overhang Bounds
10:15 – 10:45	Alessandro Arrigoni (University of L'Aquila)	Differentiation Using Multicomplex Numbers
10:45 – 11:15		Coffee Break
11:15 – 11:45	Melanie Winkler (RWTH Aachen)	Regret Minimization for Online Buffering Problems Using the Weighted Majority Algorithm
11:45 – 12:15	Rayan Chikhi (ENS Cachan Brittany)	Computational Theory for Genome Assembly
12:15 – 12:45	Charilaos Efthymiou (University of Warwick)	Deterministic Counting of Graph Colourings Using Sequences of Subgraphs

Sascha Geulen (RWTH Aachen University)
Online Capacity Maximization in Wireless Networks (Thursday, 09:15)

Abstract: In this talk we study a dynamic version of capacity maximization in the physical model of wireless communication. In our model, requests for connections between pairs of points in Euclidean space of constant dimension d arrive iteratively over time. When a new request arrives, an online algorithm needs to decide whether or not to accept the request and to assign one out of k channels and a transmission power to the channel. Accepted requests must satisfy constraints on the signal-to-interference-plus-noise (SINR) ratio. The objective is to maximize the number of accepted requests.

Using competitive analysis we study algorithms using distance-based power assignments, for which the power of a request relies only on the distance between the points. Such assignments are inherently local and particularly useful in distributed settings. For request sets with spatial lengths in $[1, \Delta]$ we derive a lower bound of $\Omega(\Delta^{d/2})$ on the competitive ratio of any deterministic online algorithm using a distance-based power assignment. Our main result is a near-optimal deterministic algorithm that is $O\left(\Delta^{(d/2)+\varepsilon}\right)$-competitive, for any constant $\varepsilon > 0$.

Harald Räcke (University of Warwick)
Oblivious Routing in the L$_p$-norm (Thursday, 09:45)

Abstract: Gupta et al. introduced a very general multi-commodity flow problem in which the cost of a given flow solution on a graph $G = (V, E)$ is calculated by first computing the link loads via a load-function l, that describes the load of a link as a function of the flow traversing the link, and then aggregating the individual link loads into a single number via an aggregation function.

We show the existence of an oblivious routing scheme with competitive ratio $O(\log n)$ and a lower bound of $\Omega(\log n / \log \log n)$ for this model when the aggregation function agg is an L$_p$-norm.

Our results can also be viewed as a generalization of the work on approximating metrics by a distribution over dominating tree metrics and the work on minimum congestion oblivious. We provide a convex combination of trees such that routing according to the tree distribution approximately minimizes the L$_p$-norm of the link loads. The embedding techniques of Bartal and Fakcharoenphol et al. can be viewed as solving this problem in the L$_1$-norm while the result on congestion minmizing oblivious routing solves it for L$_\infty$. We give a single proof that shows the existence of a good tree-based oblivious routing for any L$_p$-norm.

Rahul Savani (University of Liverpool)
The Complexity of the Homotopy Method, Equilibrium Selection, and Lemke-Howson Solutions (Thursday, 10:15)

Abstract: We show that the widely used homotopy method for solving fixpoint problems, as well as the Harsanyi-Selten equilibrium selection process for games, are PSPACE-complete to implement. A key application of our techniques yields the result that it is also PSPACE-complete to compute any of the equilibria that could be found via the classical Lemke-Howson algorithm.

Joint work with Paul W. Goldberg and Christos H. Papadimitriou.

Tobias Brunsch (Maastricht University)
Clustering - How Bad Is The k-Means++ Method? (Thursday, 11:15)

Abstract: The objective of the general clustering problem is, given a set of objects, to group those with "similar properties" into so called "clusters". A common way to rate the quality of a clustering is to use a cost function based on some kind of distance measure, and one representative of this type of clustering problems is "k-Means". A popular local search heuristic for k-means is "Lloyd's algorithm", usually called the "k-means method". The algorithm's initialization has a dramatic impact on the output quality, and recently Arthur and Vassilvitskii proposed "k-means++", an initialization technique which guarantees an expected approximation factor of $O(\log(k))$, where k denotes the number of clusters.

Though they also showed that there are instances where k-means++ does not achieve a better approximation factor in expectation, the question arose whether k-means++ yields a constant factor approximation with non-negligible (or even constant) probability. We give a negative answer to this question by constructing a class of instances on which k-means++ yields an $\Omega(\log(k))$-approximation with probability exponentially close to one.

Martin Gairing (University of Liverpool)
Covering Games: Approximation Through Non-Cooperation (Thursday, 11:45)

Abstract: We propose approximation algorithms under game-theoretic considerations. We introduce and study the general covering problem which is a natural generalization of the well-studied max-n-cover problem. In the general covering problem, we are given a universal set of weighted elements E and n collections of subsets of the elements. The task is to choose one subset from each collection such that the total weight of their union is as large as possible. In our game-theoretic setting, the choice in each collection is made by an independent player. For covering an element, the players receive a payoff defined by a non-increasing utility sharing function. This function defines the fraction that each covering player receives from the weight of the elements.

We show how to construct a utility sharing function such that every Nash Equilibrium approximates the optimal solution by a factor of $1 - 1/e$. We also prove that any sequence of unilateral improving steps is polynomially bounded. This gives rise to a polynomial-time local search approximation algorithm whose approximation ratio is best possible.

Rajiv Raman (University of Warwick)
Maximum Feasible Subsystems (Thursday, 12:15)

Abstract: In the maximum feasible subsystem problem we are given an infeasible linear system $AxRb$, (where R is either '=','$\leq$', or '$<$') and the goal is to find a solution x that satisfies the maximum number of inequalities simultaneously. This fundamental problem arises in many applications such as operations research, machine learning, computational geometry, etc. We consider the following version of the problem: All entries of the matrix A are 0 or 1, and we are given two vectors l and u. The objective is to find a non-negative vector x that simultaneously satisfies the largest subset of the inequalities $l \leq Ax \leq u$. We present hardness of approximation, as well as approximation algorithms for this problem as well as the case where the matrix A is an interval matrix, and give some applications.

Heiko Röglin (Maastricht University)
Smoothed Analysis of Multiobjective Optimization (Friday, 09:15)

Abstract: We prove that the number of Pareto-optimal solutions in any multiobjective binary optimization problem with a finite number of linear objective functions is polynomial in the model of smoothed analysis. Moreover, we give polynomial bounds on all finite moments of the number of Pareto-optimal solutions, which yields the first non-trivial concentration bound for this quantity.

This talk is based on joint work with Shang-Hua Teng.

Mike Paterson (University of Warwick)
Overhang Bounds (Friday, 09:45)

Abstract: How far can a stack of n identical blocks be made to hang over the edge of a table? The question dates back to at least the middle of the 19th century and the answer to it was widely believed to be of order $\log n$. Recently, we (Paterson and Zwick) constructed n-block stacks with overhangs of order $n^{1/3}$, exponentially better than previously thought possible. Subsequently, we (Paterson, Peres, Thorup, Winkler and Zwick) showed that order $n^{1/3}$ is best possible, resolving the long-standing overhang problem up to a constant factor.

I shall review the construction and outline the upper bound proof, which illustrates how methods founded in algorithmic complexity can be applied to a discrete optimization problem that has puzzled some mathematicians and physicists for more than 150 years.

Alessandro Arrigoni (University of L'Aquila)
Differentiation Using Multicomplex Numbers (Friday, 10:15)

Abstract: The multicomplex method for calculating derivatives and its use in numerical algorithms, i.e. optimization algorithms, is presented. A general procedure for the implementation is also described through the use of a non interpretative language, Fortran 90. In the examples the multicomplex method is compared with finite difference method, Automatic Differentiation method , i.e. Tapenade, and analytical method in terms of computational advantages and relative computational time and is shown to have a very good accuracy respect to the other methods.

Melanie Winkler (RWTH Aachen University)
Regret Minimization for Online Buffering Problems Using the Weighted Majority Algorithm (Friday, 11:15)

Abstract: Suppose a decision maker has to purchase a commodity over time with varying prices and demands. In particular, the price per unit might depend on the amount purchased and this price function might vary from step to step. The decision maker has a buffer of bounded size for storing units of the commodity that can be used to satisfy demands at later points in time. We seek for an algorithm deciding at which time to buy which amount of the commodity so as to minimize the cost. This kind of problem arises in many technological and economical settings like, e.g., battery management in hybrid cars and economical caching policies for mobile devices. A simplified but illustrative example is a frugal car driver thinking about at which occasion to buy which amount of gasoline. Within a regret analysis, we assume that the decision maker can observe the performance of a set of expert strategies over time and synthesizes the observed strategies into a new online algorithm. In particular, we investigate the external regret obtained by the well-known Randomized Weighted Majority algorithm applied to our problem. We show that this algorithm does not achieve a reasonable regret bound if its random choices are independent from step to step, that is, the regret for T steps is $\Omega(T)$. However, one can achieve regret $O(\sqrt{T})$ when introducing dependencies in order to reduce the number of changes between the chosen experts. If the price functions satisfy a convexity condition then one can even derive a deterministic variant of this algorithm achieving regret $O(\sqrt{T})$.

Our more detailed bounds on the regret depend on the buffer size and the number of available experts. The upper bounds are complemented by a matching lower bound on the best possible external regret.

Rayan Chikhi (ENS Cachan Brittany)
Computational Theory for Genome Assembly (Friday, 11:45)

Abstract: In bioinformatics, genome assembly is the computational step that consists in reconstructing a genome sequence given only a set of subsequences. This problem has attracted a lot of attention due to the genomic sequencing efforts from the past two decades. Recently, next-generation sequencing technologies are producing enormous amounts of short subsequences, hence instances of the assembly problem are even more computationally challenging. For instance, the most efficient algorithm to date requires two days on a supercomputer (40 cores, 140 Gb memory) to assemble the human genome (3 Gb). Genome assembly has several formulations as NP-hard computational problems. We survey the most popular graph-theoretic models for these formulations and discuss their complexity. We shortly describe thet heuristics used in practical implementations. A new unpublished graph model that takes into account a structural feature of the data is presented.

Charilaos Efthymiou (University of Warwick)
Deterministic Counting of Graph Colourings Using Sequences of Subgraphs (Friday, 12:15)

Abstract: In this paper we propose a deterministic algorithm for approximate counting the k-colourings of a graph G. We consider two categories of graphs. The first one is the sparse random graphs $G_{n,d/n}$, where each edge is chosen independently with probability d/n and d is fixed. The second one is a family we call *locally* α-*dense graphs* of bounded maximum degree Δ. A graph $G = (V, E)$ in this family has following property: For all $\{u, v\} \in E$ the number of vertices which are adjacent to v but not adjacent to u are at most $(1 - \alpha)\Delta$, where $\alpha \in [0, 1]$ is a parameter of the model. Our algorithm computes in polynomial time a $(1 \pm n^{-\Omega(1)})$-approximation of the logarithm of the number of k-colourings of $G_{n,d/n}$ for $k \geq 2.1d$ with high probability. Also, for G a locally α-dense graph of *bounded* maximum degree Δ it computes in polynomial time a $(1 \pm (\log n)^{-\Omega(1)})$-approximation of the logarithm of the number of k-colourings, for $k > (2 - a)\Delta$. Restricting the treewidth of the neighbourhoods in G we get improved approximation. Our algorithm is related to the algorithms of A. Bandyopadhyay et al. in SODA '06, and A. Montanari et al. in SODA '06. However, we use correlation decay properties for the *Gibbs distribution* in a completely different manner. I.e. given the graph $G = (V, E)$, we alter the graph structure in some specific region $\Lambda \subseteq V$ (by deleting edges between vertices of Λ) and then we show that the effect of this change on the marginals of Gibbs distribution, diminishes as we move away from Λ. Our approach is novel and suggests a new context for the study of deterministic counting algorithms.