FND 2012 Abstracts

Mon 9:00	**Adaptive Approximation Algorithms** R Ravi	When describing constant factor approximation algorithms to my integer programming colleagues, I get an incredulous look about the output of my algorithms having over 100 percent error, while IP techniques usually come within a few percent of optimal. I often wondered if there is a class of problems for which I would be able to do better. Adaptive approximation algorithms are algorithms for problems with multiple stages of decision making, such as multi-armed bandit problems. Here the adaptivity gap is high and hence adaptivity is essential to getting good quality solutions. Solving for the optimal adaptive policy using an integer program usually turns out to be computationally expensive while LP relaxations are still good for rounding for some classes of multi-armed bandit problems. I will describe these problems and adaptive approximation algorithms for them in this talk. Joint work wtih Anupam Gupta, Ravishankar Krishnaswamy and Marco Molinaro that appeared in FOCS 2011. Full paper at http://arxiv.org/abs/1102.3749
Mon 9:45	**LP Rounding for k-Centers with Non-uniform Hard Capaticies** Marek Cygan	In this paper we consider a generalization of the classical k-center problem with capacities. Our goal is to select k centers in a graph, and assign each node to a nearby center, so that we respect the capacity constraints on centers. The objective is to minimize the maximum distance a node has to travel to get to its assigned center. This problem is NP-hard, even when centers have no capacity restrictions and optimal factor 2-approximation algorithms are known. With capacities, when all centers have identical capacities, a 6-approximation is known with no better lower bounds than for the infinite capacity version. While many generalizations and variations of this problem have been studied extensively, no progress was made on the capacitated version for a general capacity function. We develop the first constant factor approximation algorithm for this problem. Our algorithm uses an LP rounding approach to solve this problem, and works for the case of non-uniform hard capacities, when multiple copies of a node may not be chosen and can be extended to the case when there is a hard bound on the number of copies of a node that may be selected. In addition we establish a lower bound on the integrality gap of 7(5) for non-uniform (uniform) hard capacities. Finally, for non-uniform soft capacities we present a much simpler 11-approximation algorithm, which we find as one more evidence that hard capacities are much harder to deal with. This is a joint work with MohammadTaghi Hajiaghayi and Samir Khuller.
Mon 11:00	**Prophet-Inequality Setting and Stochastic Generalized Assignment** Mohammad T. Hajiaghayi	We study the problem of online stochastic budgeted matching in bipartite graphs. There has been a recent trend of work improving over the bound of $1 - \frac{1}{e}$ most of them motivated by problems related to "Online Ad Allocation" and almost all of them consider a 1-1 matching. In this paper, we consider a weighted/budgeted model in which each edge has a weight (i.e., bid of each bidder for each query) and each vertex on the fixed side of the graph has a budget (i.e., total budget of each bidder). The weight of a matching is the minimum of the budget of each vertex and the total weight of edges matched to it when summed over all vertices. We consider this model in a more realistic stochastic setting where we know the distribution of the incoming queries in advance. We show that if the bid to the budget ratio of every bidder is at most $\frac{1}{k}$ then the natural randomized online algorithm which is based on solving a fractional LP for the expected instance has an approximation ratio of $1 - \frac{k^k}{e^k k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}$ compared to the optimal offline. We also consider a model in which the bidders do not have a budget. Each bidder may have a capacity (i.e., an upper bound on the number of ads that can be allocated to him). We present an algorithm with a tight approximation ratio of $\frac{1}{2}$ for this setting which generalizes the classic Prophet Inequality and has applications in "Display Ad Auctions" and "Ad-Cell Auctions". Time permitting, we mention a generalization of our framework for the stochastic generalized assignment problem via a continuous version of sand/barrier theorem.

Mon 11:45	**The power of recourse for online MST and TSP** Andreas Wiese	We consider online versions of MST and TSP problems with recourse. Assume that vertices of a complete metric graph appear one by one, and must be connected by a tree (respectively tour) of low cost. In the standard online setting, where decisions are irrevocable, the competitive factor of any algorithm is $\Omega(\log n)$. In our model, recourse is allowed by granting a limited number of edge rearrangements per iteration. More than 20 years ago, Imase and Waxman (1991) conjectured that constant-competitive solutions can be achieved with a constant (amortized) number of rearrangements. In this talk I will present an algorithm that solves this conjecture for MSTs in the amortized setting. Unlike in offline TSP variants, the standard double-tree and shortcutting approach does not give constant guarantees in the online setting. However, a non-trivial robust shortcutting technique allows to convert trees into tours at the loss of small factors, implying the conjecture of Imase and Waxman for tours. For the non-amortized setting, we show that if the whole sequence is known in advance then already with arbitrarily few rearrangements we can construct robust sequences of spanning trees which are constant-competitive at each iteration.
Mon 13:30	**Algorithmic Applications of Baur-Strassen's Theorem: Shortest Cycles, Diameter and Matchings** Piotr Sankowski	Consider a directed or an undirected graph with integral edge weights from the set $[-W, W]$, that does not contain negative weight cycles. In this paper, we introduce a general framework for solving problems on such graphs using matrix multiplication. The framework is based on the usage of Baur-Strassen's theorem and of Strojohann's determinant algorithm. It allows us to give new and simple solutions to the following problems: - Finding Shortest Cycles - We give a simple $O(Wn^{\omega})$ time algorithm for finding shortest cycles in undirected and directed graphs. For directed graphs this matches the time bounds obtained in 2011 by Roditty and Vassilevska-Williams. On the other hand, no algorithm working in $\tilde{O}(Wn^{\omega})$ time was previously known for undirected graphs with negative weights. - Computing Diameter - We give a simple $O(Wn^{\omega})$ time algorithm for computing a diameter of an undirected or directed graphs. This considerably improves the bounds of Yuster from 2010, who was able to obtain this time bound only in the case of directed graphs with positive weights. In contrary, our algorithm works in the same time bound for both directed and undirected graphs with negative weights. - Finding Minimum Weight Perfect Matchings - We present an $O(Wn^{\omega})$ time algorithm for finding minimum weight perfect matchings in undirected graphs. This resolves an open problem posted by Sankowski in 2006, who presented such an algorithm but only in the case of bipartite graphs. We believe that the presented framework can find applications for solving larger spectra of related problems. As an illustrative example we apply it to the problem of computing a set of vertices that lie on cycles of length at most t, for some t. We give a simple $O(Wn^{\omega})$ time algorithm for this problem that improves over the $O(tWn^{\omega})$ time algorithm given by Yuster in 2011. This is joint work with Marek Cygan and Hal Gabow.
Mon 14:15	**Faster Deterministic Fully-Dynamic Graph Connectivity** Christian Wulff-Nilsen	We give new deterministic bounds for fully-dynamic graph connectivity. Our data structure supports updates (edge insertions/deletions) in $O(\frac{\log^2 n}{\log\log n})$ amortized time and connectivity queries in $O(\frac{\log n}{\log\log n})$ worst-case time, where n is the number of vertices of the graph. This improves the deterministic data structures of Holm, de Lichtenberg, and Thorup (STOC 1998, J.ACM 2001) and Thorup (STOC 2000) which both have $O(\log^2 n)$ amortized update time and $O(\frac{\log n}{\log\log n})$ worst-case query time.
Mon 15:30	**Directed Steiner Tree and the Lasserre Hierarchy** Thomas Rothvoss	The goal for the Directed Steiner Tree problem is to find a minimum cost tree in a directed graph $G = (V, E)$ that connects all terminals X to a given root r. It is well known that modulo a logarithmic factor it suffices to consider acyclic graphs where the nodes are arranged in $L \leq \log \lvert X \rvert$ levels. Unfortunately the natural LP formulation has a $\lvert X \rvert^{\frac{1}{2}}$ integrality gap already for five levels. We show that for every L, the $O(L)$-round Lasserre strengthening of this LP has integrality gap $O(L \log \lvert X \rvert)$. This provides a polynomial time $\lvert X \rvert^{\epsilon}$-approximation and a $O(\log^3 \lvert X \rvert)$ approximation in $O(n^{\log \lvert X \rvert})$ time, matching the best known approximation guarantee obtained by a greedy algorithm of Charikar et al. The paper is available under http://arxiv.org/abs/1111.5473.

Mon 16:15	**Improved Distance Sensitivity Oracles via Fast Single-Source Replacement Paths** Fabrizio Grandoni	A *distance sensitivity oracle* is a data structure which, given two nodes s and t in a directed edge-weighted graph G and an edge e, returns the shortest length of an s-t path not containing e, a so called *replacement path* for the triple (s, t, e). Such oracles are used to quickly recover from edge failures. In this talk we consider the case of integer weights in the interval $[-M, M]$, and present the first distance sensitivity oracle that achieves simultaneously subcubic preprocessing time and sublinear query time. More precisely, for a given parameter $\alpha \in [0, 1]$, our oracle has preprocessing time $\tilde{O}(Mn^{\omega+\frac{1}{2}} + Mn^{\omega+\alpha(4-\omega)})$ and query time $\tilde{O}(n^{1-\alpha})$. Here $\omega < 2.373$ denotes the matrix multiplication exponent. For a comparison, the previous best oracle for small integer weights has $\tilde{O}(Mn^{\omega+1-\alpha})$ preprocessing time and (superlinear) $\tilde{O}(n^{1+\alpha})$ query time [Weimann,Yuster-FOCS'10]. The main novelty in our approach is an algorithm to compute all the replacement paths from a given source s, an interesting problem on its own. We can solve the latter *single-source replacement paths* problem in $\tilde{O}(APSP(n, M)))$ time, where $APSP(n, M) < \tilde{O}(M^{0.681}n^{2.575})$ [Zwick-JACM'02] is the runtime for computing all-pairs shortest paths in a graph with n vertices and integer edge weights in $[-M, M]$. For positive weights the runtime of our algorithm reduces to $\tilde{O}(Mn^{\omega})$. This matches the best known runtime for the simpler *replacement paths* problem in which both the source s and the target t are fixed [Vassilevska-SODA'11].
Tue 9:00	**Improving on Christofides' algorithm for the s-t path TSP** David Shmoys	We present an approximation algorithm for the s-t path traveling salesman problem (TSP) for an arbitrary metric, which is guaranteed to find a solution of cost within a factor of the golden ratio of optimal in polynomial time. In this problem, we are given pairwise distances among n points that satisfy the triangle inequality, and two pre-specified endpoints s and t, and the problem is to find a shortest Hamiltonian path between s and t; Hoogeveen showed that the natural variant of the classic TSP algorithm of Christofides is a $\frac{5}{3}$-approximation algorithm for this problem, and this asymptotically tight bound has been the best approximation ratio known until now. We modify this algorithm so that it chooses the initial spanning tree based on an optimal solution to a natural linear programming relaxation, rather than a minimum spanning tree; we prove this simple but crucial modification leads to an improved approximation ratio, surpassing the 20-year-old barrier set by the natural Christofides' algorithm variant. This is joint work with Hyung-Chan An and Robert Kleinberg.
Tue 9:45	**Reachability in Graph Timelines** Jakub Łącki	We consider the problem of maintaining information about graphs with history - so called graph timeline. In this problem we are given a sequence of graphs $G_1, ..., G_t$, where two consecutive graphs in the sequence differ by addition or removal of one edge. We aim to devise algorithms that require as little preprocessing and memory as possible and are able to efficiently answer queries about the existence of paths in the graph timeline in a given time interval. This is joint work with Piotr Sankowski.
Tue 11:00	**Upward Max Min Fairness** Haim Kaplan	Often one would like to allocate shared resources in a fair way. A common and well studied notion of fairness is *Max-Min Fairness*, where we first maximize the smallest allocation, and subject to that the second smallest, and so on. We consider a networking application where multiple commodities compete over the capacity of a network. In our setting each commodity has multiple possible paths to route its demand (for example, a network using MPLS tunneling). In this setting, the only known way of finding a max-min fair allocation requires an iterative solution of multiple linear programs. Such an approach, although polynomial time, scales badly with the size of the network, the number of demands, and the number of paths. More importantly, a network operator has limited control and understanding of the inner working of the algorithm. Finally, this approach is inherently centralized and cannot be implemented via a distributed protocol. In this work we introduce Upward Max-Min Fairness, a novel relaxation of Max-Min Fairness and present a family of simple dynamics that converge to it. Moreover, we present an efficient combinatorial algorithm for finding an upward max-min fair allocation which is a natural extension of the well known Water Filling Algorithm for a multiple path setting. We test the expected behavior of this new algorithm and show that on realistic networks upward max-min fair allocations are comparable to the max-min fair allocations both in fairness and in network utilization.

Tue 11:45	**Edge-connectivity augmentation** Attila Bernáth	This talk wants to give an overview about results on edge-connectivity augmentation. We will mainly consider the problem of satisfying the given edge-connectivity requirement by adding a minimum number of undirected graph edges, however the starting network can even be directed in some cases. The edge-connectivity requirement can have many forms. It is either given as a **global** parameter (we want to get a k-edge-connected graph), or it is given for every pair of nodes (this is called **local** edge-connectivity augmentation). We will also consider the setting of the **node-to-area** connectivity augmentation: here we are given some subsets (called *areas*), and a positive integer $r(W)$ for every area W, and we want to add new edges so that the local edge-connectivity between any area W and any node outside W is at least $r(W)$. We will discuss the common abstract framework for all these problems, the framework of **covering a skew-supermodular function** by graph edges, and the most widely used technique in this context, the **splitting-off** technique.
Tue 15:30	**Improved LP-rounding approximation algorithm for k-level uncapacitated facility location** Jarek Byrka	We study the k-level uncapacitated facility location problem, where clients need to be connected with paths crossing open facilities of k types (levels). In this talk I will present an approximation algorithm that for any constant k, in polynomial time, delivers solutions of cost at most α_k times OPT , where α_k is an increasing function of k, with $\lim_{k\to\infty} \alpha_k = 3$. Our algorithm rounds a fractional solution to an extended LP formula- tion of the problem. The rounding builds upon the technique of itera- tively rounding fractional solutions on trees (Garg, Konjevod, and Ravi SODA'98) originally used for the group Steiner tree problem. We improve the approximation ratio for k-UFL for all $k \geq 3$, in particular we obtain the ratio equal 2.02, 2.14, and 2.24 for $k = 3, 4$, and 5.
Tue 16:15	**A Tight Linear Time $\frac{1}{2}$-Approximation for Unconstrained Submodular Maximization** Seffi Naor	We consider the Unconstrained Submodular Maximization problem in which we are given a non-negative submodular function f, and the objective is to find a subset S maximizing $f(S)$. This is considered to be one of the most basic submodular optimization problems, having a wide range of applications, both in theory and in more practical settings. Some well known problems captured by Unconstrained Submodular Maximization include Max-Cut, Max-DiCut, and variants of Max-SAT and maximum facility location. We present a simple randomized linear time algorithm achieving a tight approximation guarantee of $\frac{1}{2}$, thus matching the known hardness result of Feige et al. Our algorithm is based on an adaptation of the greedy approach which exploits certain symmetry properties of the problem.
Wed 9:00	**Approximation algorithms for min-cost (S,T)-connectivity** Joseph Cheriyan	One of the well-known problems in network design is to find a minimum-cost strongly-connected spanning subgraph of a directed network. We study a natural generalization on directed networks called the min-cost (S,T) connectivity problem. The talk will introduce this problem and summarize some of the algorithmic results. The paper (with Bundit Laekhanukit) is on my webpage: http://www.math.uwaterloo.ca/~jcheriyan
Wed 9:45	**On the diameter of network polytopes** Leen Stougie	The Hirsch Conjecture was formulated by Warren M. Hirsch in 1957 and states that a n-dimensional polyhedron with m facets has combinatorial diameter at most $m - n$. The conjecture was first disproven for unbounded polyhedra and remained unsettled for polytopes until Santos' breakthrough in 2010, showing that it is also false for polytopes in general. It is still a fundamental question whether the diameter of a polyhedron can be bounded up by a polynomial in n and m. In this lecture I will present some results and open problems on the diameter of network polytopes: polytopes related to optimisation problems on graphs. For some classes of network polytopes the Hirsch Conjecture has been shown to hold and for others main progress has been made. I will give an example of a typical result in this class and sho the typical type of proof technique used to obtain results. The lecture will be concluded with an overview of results and some challenging open problems.
Wed 11:00	**Strongly Truthful Mechanisms** Amos Fiat	We present the notion of strongly truthful mechanisms and give applications thereof. These include externality resistent mechanisms.

Wed 11:45	**How to use complex semidefinite programming for Max-k-Cut** Alantha Newman	A little more than a decade ago, Goemans and Williamson showed how to formulate and round a complex semidefinite program (CSP) to give the best-known approximation guarantee for the Max-3-Cut problem. They left open the problem of how to apply their techniques to the Max-k-Cut for general k. They point out that it does not seem staightforward or even possible to formulate a good quality complex semidefinite program for the general problem, which presents a barrier to the application of their techniques. While other works (e.g by Frieze and Jerrum) have analyzed other semidefinite programming based algorithms for the Max-k-Cut problem and these algorithms are quite simple, their analyses are quite complicated. We will discuss how the tools from Goemans and Williamson's CSP paper can be applied to the Max-k-Cut problem for general k, which leads to weaker approximation guarantees but to an arguably simpler analysis.
Thu 9:00	**Optimal Online Buffer Scheduling for Block Devices** Matthias Englert	We introduce a buffer scheduling problem for block operation devices in an online setting. We consider a stream of items of different types to be processed by a block device. The block device can process all items of the same type in a single step. To improve the performance of the system a buffer of size k is used to store items in order to reduce the number of operations required. Whenever the buffer becomes full a buffer scheduling strategy has to select one type and then a block operation on all elements with this type that are currently in the buffer is performed. The goal is to design a scheduling strategy that minimizes the number of block operations required. In this paper we consider the online version of this problem, where the buffer scheduling strategy must make decisions without knowing the future items that appear in the input stream. Our main result is the design of an O(loglog k)-competitive online randomized buffer scheduling strategy. The bound is asymptotically tight. As a byproduct of our LP-based techniques, we obtain a randomized offline algorithm that approximates the optimal number of block-operations to within a constant factor. (joint work with Anna Adamaszek, Artur Czumaj, and Harald Raecke)
Thu 9:45	**Capacitated Network Design** Deeparnab Chakrabarty	I am going to talk about the capacitated network design problem which asks to find the cheapest subgraph of a *capacitated* graph which satisfies certain connectivity conditions. When all capacities are same, the problem boils down to the decently well understood SNDP for which Jain gave his beautiful 2-approximation. One probably needs to dig deeper to extract beauty in the case of general capacities; this talk will try and sketch what's known about the problem.
Thu 11:00	**Robust Expansion** Kunal Talwar	I will talk about the problem of robust expansion. Given a bipartite graph $G = (U, V, E)$, the robust expansion of a set $A \subseteq U$ is the ratio $\frac{\lvert B\rvert}{\lvert A\rvert}$, where $B \subseteq V$ is the smallest set such that $E(A, B) \geq \alpha E(A, V)$. In words, we want to capture an α fraction of the edges leaving A, and B is the smallest set that does that. Note that when alpha is 1, this is just the vertex expansion of the set A. The α-robust expansion of the graph is then defined to be the minimum over all A of the robust expansion of A. I will talk about the problem of approximating the robust expansion of a regular bipartite graph. (Joint work with Aditya Bhaskara, Uri Feige and Ravishankar Krishnaswamy)
Thu 11:45	**Linear-Programming based Techniques for Minimum Latency Problems** Chaitanya Swamy	We introduce a problem that is a common generalization of the uncapacitated facility location (UFL) and minimum latency (ML) problems. In our model, as in UFL, there are facilities and clients, and we need to open facilities and assign clients to facilities. But in addition we need to find a tour visiting the facilities that determines the time by which a facility will be activated or replenished. The goal is to minimize the facility-opening costs and the delays incurred by the clients, where the latter term takes into account both the time taken for a client to reach its assigned facility and the time taken to activate this facility. Besides being a natural problem of interest, this problem unifies various diverse combinatorial-optimization problems, and our results yield various insights for some of these problems as well. We give a polylogarithmic approximation algorithm for this problem, which is tight in that any improvement on this would yield an analogous improvement for the group Steiner tree problem, which is long-standing open problem. Our techniques yield LP-based methods for minimum-latency (and other vehicle routing) problems. Specifically, we obtain two LP-relaxations for ML, which we believe offer a promising source for improving the approximation for ML. We can prove an upper bound on the integrality gap by drawing upon various ideas in scheduling, facility location and polyhedral theory.

Thu 15:30	**The Vital Core Connectivity Problem** Sylvia Boyd	Let $G = (V, E)$ be an edge-weighted complete graph representing a network in which the vertices are partitioned into two classes–vital vertices V^*, which represent the vital core of the network, and secondary vertices $V \setminus V^*$. We consider the vital core connectivity problem $\mathrm{VCC}(k, r, V^*)$, which is the problem of finding a minimum weight spanning multi-subgraph of G which is k-edge connected overall and whose vital core remains at least r-edge connected even if some or all of the secondary vertices are removed. This problem arises naturally in many practical applications in which one wishes to design a network at minimum cost which will not only survive the loss of a certain number of links overall, but for which the vital core remains at least r-edge connected even if some or all of the secondary centres are lost. We give a p-approximation algorithm for this problem, where $\frac{3}{2} \leq p \leq \frac{19}{6}$, depending on the values of k and r. We show that our algorithm has a better performance guarantee than can be obtained by combining other known heuristics for related problems.
Thu 16:15	**$(1+\epsilon)$-Approximation for Facility Location in Data Streams** Artur Czumaj	We consider the Euclidean facility location problem with uniform opening cost, and present two main results: - a very simple and fast PTAS with the running time $O(n \log^2 n \log \log n)$ for constant ϵ, - the first $(1+\epsilon)$-approximation algorithm for the cost of the facility location problem for dynamic geometric data streams, using polylogarithmic space. The key, novel ingredient of our algorithms is a simple decomposition scheme that enables to partition the input points into disjoint point sets that can be then considered independently.
Fri 9:00	**An $O(\log k)$-competitive Algorithm for Generalized Caching** Harald Räcke	In the generalized caching problem, we have a set of pages and a cache of size k. Each page p has a size $w_p \geq 1$ and fetching cost c_p for loading the page into the cache. At any point in time, the sum of the sizes of the pages stored in the cache cannot exceed k. The input consists of a sequence of page requests. If a page is not present in the cache at the time it is requested, it has to be loaded into the cache incurring a cost of c_p. We give a randomized $O(\log k)$-competitive online algorithm for the generalized caching problem, improving the previous bound of $O(\log^2 k)$ by Bansal, Buchbinder, and Naor. This improved bound is tight and of the same order as the known bounds for the classic problem with uniform weights and sizes. We use the same LP based techniques as Bansal et al. but provide improved and slightly simplified methods for rounding fractional solutions online. This is joint work with Anna Adamaszek, Artur Czumaj, and Matthias Englert.
Fri 9:45	**Claw-Free k-in-a-Tree** Erik Jan van Leeuwen	The k-in-a-Tree problem asks to find an induced subgraph that contains k given terminals and that is a tree. This problem is NP-hard in general. The closely related k-in-a-Cycle and k-in-a-Path problems are even harder, as they are NP-hard when $k = 2$ and $k = 3$ respectively. We consider these problems on $K_{1,3}$-free (i.e. claw-free) graphs, and show that they are all fixed-parameter tractable in k on such graphs. In fact, we show that the more general Induced Disjoint Paths problem is fixed-parameter tractable on claw-free graphs. In contrast, we observe that Induced Disjoint Paths does not have a polynomial kernel on claw-free graphs, and that all previously mentioned problems are NP-complete on $K_{1,4}$-free graphs even for constant k. Based on joint work with Daniel Paulusma and Petr A.G̃olovach that will appear in ESA 2012.