

ISSN 2186-7437

NII Shonan Meeting Report

No. 2016-7

Current Trends in Combinatorial Optimization

Takuro Fukunaga
László Végh
R. Ravi

April 11–14, 2016



National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Current Trends in Combinatorial Optimization

Organizers:

Takuro Fukunaga (National Institute of Informatics)

R. Ravi (Carnegie Mellon University)

László Végh (London School of Economics)

April 11–14, 2016

1 Overview

Combinatorial optimization is a field in the intersection of applied discrete mathematics, operations research, and computer science. The goal in the combinatorial optimization is design of efficient algorithms for computing a good quality solution to problems on discrete structures whose solution space is subject to combinatorial explosion. It has numerous applications in various fields of science, and is regarded as a fundamental technological advance in processing data for optimal solutions and structures. Clearly its importance is increasing with the growing applications of data science, and these algorithmic methods are also gathering much public attention nowadays in fields as diverse as online dating to package routing.

The theory of combinatorial optimization has recorded remarkable developments in the past fifty years. However, many fundamental questions still remain unsolved. An example of such questions is the one related to the traveling salesperson problem (TSP). This is the problem of computing a shortest tour visiting all the cities in a map, where the distance between pair of cities is given as input. Since the number of tours corresponds to the number of cyclic permutations of the cities, it is an exponentially large number and hence a victim of combinatorial explosion. Nevertheless, for the TSP with distances that obey the triangle inequality, Christofides presented an elegant 1.5-approximation algorithm in 1976. This method is guaranteed to output a tour of length at most 1.5 times the shortest length of arbitrary feasible solution in time that grows with a polynomial dependence on the number of cities. One of the most popular current questions in combinatorial optimization asks whether there is a polynomial-time $4/3$ -approximation algorithm for metric TSP. In 2011, several groups made progress for a special case defined by the shortest path hop distances in undirected graphs. They gave approximation algorithms that achieve approximation ratios better than 1.5 for this special case. However there is still no known algorithm that outperforms Christofides algorithm in instances with arbitrary metrics.

Another example is the bin packing problem. In the bin packing problem, we are given a set of items whose size are specified by real numbers in $[0,1]$, and we are asked to pack the items into the smallest number of bins of size one. Several groups of researchers proved that various greedy algorithms achieve

constant approximation for the problem in 1980s. Then, in 1982, Karmarkar and Karp presented an algorithm that computes a packing using at most $\text{OPT} + O(\log^2 \text{OPT})$ bins, where OPT denotes the number of bins used by the optimal solutions. Their algorithm is based on a linear programming relaxation of the problem formulated by Gilmore and Gomory in 1961. Their analysis is built on a deep understanding on the structure of the relaxation. Since their work, there was no significant progress for three decades although it is conjectured that the bin packing problem has an algorithm that computes a packing with $\text{OPT}+1$ bins. In 2013, Rothvo improved the algorithm of Karmarkar and Karp. His algorithm computes a solution with $\text{OPT} + O(\log \text{OPT} \log \log \text{OPT})$. He improved the rounding phase of method of Karmarkar and Karp using recently developed algorithmic tools from discrepancy theory.

Besides these examples, we have numerous combinatorial optimization problems for which technical breakthrough are expected. To advance development of combinatorial optimization further, wide collaboration among researchers is necessary. The purpose of this meeting is to stimulate communication among researchers who are active in this field. To enable this, we plan to invite leading and active researchers, and provide an opportunity to share the most recent information on the new findings in this field.

2 Schedule

We had 22 long talks and 8 short talks. 30 minutes are allocated for each long talk, and 15 minutes are allocated for each short talk.

April 11 (Mon)

9:00–11:00 Introductions

11:00–12:00 2 long talks

- László Végh: Rescaled coordinate descent methods for linear programming
- R. Ravi: Sending secrets swiftly

14:00–15:00 Introductions

16:00–18:00 4 long talks

- Alina Ene: Routing under balance
- Debmalya Panigrahi: Online algorithms for multi-commodity network design
- Tobias Mömke: Airports and railways: facility location meets network design
- Kristóf Bérczi: Degree-sequences of highly-connected digraphs

April 12 (Tue)

9:00–10:30 2 long talks + 2 short talks

- Alantha Newman: The alternating stock size problem and the gasoline puzzle (long)

- Yutaro Yamaguchi: How to make a bipartite graph DM-irreducible by adding edges (long)
- Tamás Király: Minimizing submodular functions on diamonds via generalized fractional matroid matchings (short)
- Neil Olver: On the integrality gap of the cut LP for prize-collecting Steiner forest (short)

11:00–12:00 2 long talks

- Friedrich Eisenbrand: Max-sum diversity via convex programming
- Tasuku Soma: Non-convex compressed sensing with the sum-of-squares method

14:00–15:00 2 long talks

- András Sebő: The salesman’s improved paths
- Anupam Gupta: The independent set problem on sparse graphs

16:00–17:30 2 long talks + 2 short

- Seeun William Umboh: Robust algorithms for noisy minor-free and bounded treewidth graphs (long)
- Bundit Laekhanukit: Approximating directed Steiner problems via tree embedding (long)
- Yasushi Kawase: The secretary problem with a choice function (short)
- Zachary Friggstad: Local search yields a PTAS for k -means in doubling metrics (short)

April 13 (Wed)

9:00–10:30 2 long talks + 2 short talks

- Gyula Pap: Blocking arborescences (long)
- Karthekeyan Chandrasekaran: Stabilizing unstable graphs through minimum modification (long)
- Naonori Kakimura: Threshold influence model for allocating advertising budgets (short)
- Satoru Fujishige: A solution to the random assignment problem with a matroidal family of goods (short)

11:00–12:00 2 long talks

- Seffi Naor: Multi-label classification with pairwise relations
- Sanjeev Khanna: Approximate matchings in dynamic graph streams

April 14 (Thu)

9:00–10:30 2 long talks + 2 short talks

- Jakub Tarnawski: Constant-factor approximation for ATSP with two edge weights (long)

- Venkatesan Guruswami: Coloring low-discrepancy hypergraphs and promise constraint satisfaction (long)
- Nikhil Bansal: An improve approximation for weighted completion time on unrelated machines (short)
- Hiroshi Hirai: Combinatorial algorithms for some multiflow problems and related network designs (short)

11:00–12:00 2 long talks

- Viswanath Nagarajan: Approximation algorithms for inventory problems with submodular or routing costs
- Jannik Matuschke: Recent developments in robust network flows

14:00–16:00 Wrap-up

3 Abstracts of Talks

Rescaled coordinate descent methods for linear programming

László Végh (London School of Economics)

Simple coordinate descent methods such as von Neumann’s algorithm or Perceptron, both developed in the 50s, can be used to solve linear programming feasibility problems. Their convergence rate depends on the condition measure of the problem at hand, and is typically not polynomial. Recent work of Chubanov (2012, 2014), related to prior work of Betke (2004), has gathered renewed interest in the application of these methods in order to obtain polynomial time algorithms for linear programming. We present two algorithms that fit into this line of research. Both our algorithms alternate between coordinate descent steps and rescaling steps, so that either the descent step leads to a substantial improvement in terms of the convergence, or we can infer that the problem is ill conditioned and rescale in order to improve the condition measure. In particular, both algorithms are based on the analysis of a geometrical invariant of the LP problem, used as a proxy for the condition measure, that appears to be novel in the literature. This is joint work with Daniel Dadush (CWI) and Giacomo Zambelli (LSE).

Sending secrets swiftly

R. Ravi (Carnegie Mellon University)

We study the problem of computing a minimum time schedule to spread rumors in a given graph under several models: In the radio model, all neighbors of a transmitting node listen to the messages and are able to record it only when no other neighbor is transmitting; In the wireless model (also called the edge-star model), each transmitter is at a different frequency to which any neighbor can tune to, but only one neighboring transmission can be accessed in this way; In the telephone model, the set of transmitter-receiver pairs form a matching in the graph. The rumor spreading problems assume a message at one or several nodes

of the graph that must reach a target node or set of nodes. The transmission proceeds in synchronous rounds under the rules of the corresponding model. The goal is to compute a schedule that completes in the minimum number of rounds. We'll present a comprehensive study of approximation algorithms for these problems, and show some reductions from the harder to the easier models for special demands

In the telephone model, we consider a new multi-commodity version where the rumors are no longer from a single node but from different sources and intended for specific destinations. Using ideas from earlier work, we present a very simple greedy algorithm for the basic multicast problem with logarithmic performance guarantee and adapt it to this extension to give an approximation algorithm with performance ratio $2^{O(\log \log k \sqrt{\log k})}$ for k source-sink pairs.

Joint work with Jennifer Iglesias (CMU), Afshin Nikzad (Stanford), Rajmohan Rajaraman (Northeastern) and Ravi Sundaram (Northeastern).

Routing under balance

Alina Ene (University of Warwick)

Over the last few decades, a tremendous amount of progress has been made toward understanding fundamental routing problems in graphs. This research agenda has led to ground-breaking results ranging from oblivious routing schemes to routing on disjoint paths and finding maximum flows and cuts, among many others. In particular, a very interesting picture has emerged: in undirected graphs, these problems admit very good approximation guarantees and very fast running times; in stark contrast, in directed graphs, these problems are considerably harder and strong negative results abound.

In this talk, we describe recent work that aims to bridge the gap between the undirected setting and the worst-case directed setting by providing a more fine-grained understanding of routing problems in directed graphs. We introduce the notion of balance for directed graphs, which is the maximum ratio between the total incoming and outgoing capacities across any cut; several important families of graphs are nearly balanced, in particular, Eulerian graphs and residual graphs of approximate undirected maximum flows. We illustrate the power of this notion by providing an oblivious routing scheme for single-source instances with a poly-logarithmic competitive ratio, and a nearly linear time algorithm for finding an approximate maximum s-t flow and minimum s-t cut.

Online algorithms for multi-commodity network design

Debmalya Panigrahi (Duke University)

Network design problems can be broadly classified into single-sink and multi-commodity problems, depending on whether all terminals have a common destination or not. In the offline model, standard reductions are known from the more general multi-commodity problems to their corresponding single-sink versions (Chekuri, Hajiaghayi, Kortsarz, and Salavatipour, FOCS '06, SODA '07). In the online setting, however, there has historically been a time lag between a single-sink algorithm and its multi-commodity generalization. Examples include edge-weighted Steiner tree (Imase and Waxman, SIDMA '91) and Steiner forest

(Azar, Awerbuch, and Bartal, SODA '96), node-weighted Steiner tree (Naor, Panigrahi, and Singh, FOCS '11) and Steiner forest (Hajiaghayi, Liaghat, and Panigrahi, FOCS '13), etc. In fact, the utility of a reduction gadget is higher for the online problems in the following sense: while the logarithmic loss in approximation suffered by the offline reduction gadget makes it unusable for problems with constant approximations like edge-weighted Steiner tree/forest, the online variants of these problems already have logarithmic lower bounds on their competitive ratio and can therefore utilize the gadget without significantly changing their approximation guarantees. In this talk, I will describe such an online reduction gadget using a primal-dual approach, and use it to unify previously known results in Steiner problems and obtain new ones in buy-at-bulk network design. This talk will be primarily based on joint work with Deeparnab Chakrabarty, Alina Ene, and Ravishankar Krishnaswamy.

Airports and railways: facility location meets network design

Tobias Mömke (Saarland University)

We introduce a new framework of Airport and Railway Problems, which combines capacitated facility location with network design. In this framework we are given a graph with weights on the vertices and on the edges, together with a parameter k . The vertices of the graph represent cities, and weights denote respectively the costs of opening airports in the cities and building railways that connect pairs of cities. The parameter k can be thought of as the capacity of an airport. The goal is to construct a minimum cost network of airports and railways connecting the cities, where each connected component in the network spans at most k vertices, contains an open airport, and the network satisfies some additional requirements specific to the problem in the framework.

We consider two problems in this framework. In the AR_F problem there are no additional requirements for the network. This problem is related to capacitated facility location. In the AR_P problem, we require each component to be a path with airports at both endpoints. AR_P is a relaxation of the capacitated vehicle routing problem (CVRP).

We consider the problems in the two-dimensional Euclidean setting. We show that both AR_F and AR_P are NP-hard, even for uniform vertex weights (i.e., when the cost of building an airport is the same for all cities). On the positive side, we provide polynomial time approximation schemes for AR_F and AR_P when vertex weights are uniform. We also investigate AR_F and AR_P for $k = \text{infinity}$. In this setting we present an exact polynomial time algorithm for AR_F with general vertex costs, which also works for general edge costs. In contrast to AR_F, AR_P remains NP-hard when $k = \text{infinity}$, and we present a polynomial time approximation scheme for general vertex weights.

Degree-sequences of highly-connected digraphs

Kristóf Bérczi (MTA-ELTE, Budapest)

There is an extensive literature of problems concerning degree sequences of graphs or digraphs with some prescribed properties such as simplicity or

k -connectivity. For example, Edmonds characterized the degree-sequences of simple k -edge-connected undirected graphs, while Wang and Kleitman solved the corresponding problem for simple k -node-connected graphs.

By generalizing a recent result of Hong, Liu, and Lai on characterizing the degree-sequences of simple strongly connected directed graphs, we provide a characterization for degree-sequences of simple k -node-connected digraphs. More generally, we solve the directed node-connectivity augmentation problem when the augmenting digraph is degree-specified and may only use certain edges of the starting digraph and edges not contained in it. The analogous problem for edge-connectivity augmentation is also solved in the special case when the edge-connectivity is to be increased by one.

Joint work with András Frank.

The alternating stock size problem and the gasoline puzzle

Alantha Newman (CNRS, Grenoble)

Given a set S of integers whose sum is zero, consider the problem of finding a permutation of these integers such that: (i) all prefixes of the ordering are non-negative, and (ii) the maximum value of a prefix sum is minimized. Kellerer et al. referred to this problem as the “Stock Size Problem” and showed that it can be approximated to within $3/2$. They also showed that an approximation ratio of 2 can be achieved via several simple algorithms.

We consider a related problem, which we call the “Alternating Stock Size Problem”, where the number of positive and negative integers in the input set S are equal. The problem is the same as above, but we are additionally required to alternate the positive and negative numbers in the output ordering. This problem also has several simple 2-approximations. We show that it can be approximated to within 1.79.

Then we show that this problem is closely related to an optimization version of the Gasoline Puzzle due to Lovasz, in which we want to minimize the size of the gas tank necessary to go around the track. We give a 2-approximation for this problem, based on rounding an LP relaxation whose feasible solutions are convex combinations of permutation matrices.

This is joint work with Heiko Roeglin (Universitaet Bonn) and Johanna Seif (ENS Lyon).

How to make a bipartite graph DM-irreducible by adding edges

Yutaro Yamaguchi (Osaka University)

The Dulmage-Mendelsohn decomposition (or the DM-decomposition for short) gives a unique partition of the vertex set of a bipartite graph reflecting the structure of all the maximum matchings therein. A bipartite graph is said to be DM-irreducible if the DM-decomposition consists of a single component, which is equivalent (under the connectivity) to the condition that every edge is contained in some perfect matching. It is not difficult to see that every complete bipartite graph is DM-irreducible, but completeness is not necessary for bipartite graphs to be so. Then, there arises a natural question: how efficiently is

a given bipartite graph made DM-irreducible by adding edges? Specifically, we focus on the problem of finding a minimum number of additional edges to attain the DM-irreducibility. When the input bipartite graph is balanced (i.e., the left-side and right-side vertex sets have the same cardinality) and has a perfect matching, this problem reduces to making a directed graph strongly connected by adding edges, for which the minimum number of additional edges was characterized by Eswaran and Tarjan (1976). In this talk, we provide a solution to the general case by showing how to make an arbitrary bipartite graph DM-irreducible by adding a minimum number of edges. This talk is based on a joint work with Satoru Iwata and Jun Kato (University of Tokyo).

Minimizing submodular functions on diamonds via generalized fractional matroid matchings

Tamás Király (ELTE Budapest)

In the talk I will describe a polynomial-time algorithm for the problem of minimizing submodular functions on the product of diamonds. Minimizing a submodular function on a distributive lattice is known to be reducible to standard submodular minimization; however, the complexity of the problem on arbitrary product lattices is wide open. Krokhn and Larose showed that certain lattice operations preserve the tractability of the problem, and as a corollary obtained that submodular minimization is polynomial-time solvable on the product of copies of the pentagon. These techniques do not seem to work for the diamond lattice; before the present result, only a pseudo-polynomial time algorithm was known, due to Kuivinen. Our polynomial-time algorithm is based on an interesting connection to a result of Gijswijt and Pap on the weighted fractional matroid matching problem. Joint work with Satoru Fujishige, Kazuhisa Makino, Kenjiro Takazawa, and Shin-ichi Tanigawa.

On the integrality gap of the cut LP for prize-collecting Steiner forest

Neil Olver (VU Amsterdam)

The current best approximation factor for the prize-collecting Steiner forest problem is 2.54 [Hajiaghayi-Jain 06], based on rounding the natural cut LP. Many people have asked whether an approximation factor of 2, matching the results for (non-prize collecting) Steiner forest and (LP-relative) prize-collecting Steiner tree, might be possible. We show that the cut LP has an integrality gap strictly worse than 2. The construction is motivated by an attempt to extend the approximate integer decomposition approach of Chekuri-Shepherd to the prize-collecting setting.

(Joint work with J. Koenemann, R. Ravi, G. Schaefer and C. Swamy)

Max-sum diversity via convex programming

Friedrich Eisenbrand (EPFL)

Diversity maximization is an important concept in information retrieval, computational geometry and operations research. Usually, it is a variant of the

following problem: Given a ground set, constraints, and a function $f(\cdot)$ that measures diversity of a subset, the task is to select a feasible subset S such that $f(S)$ is maximized. The *sum-dispersion* function $f(S) = \sum_{x,y \in S} d(x,y)$, which is the sum of the pairwise distances in S , is in this context a prominent diversification measure. The corresponding diversity maximization is the *max-sum* or *sum-sum diversification*. Many recent results deal with the design of constant-factor approximation algorithms of diversification problems involving sum-dispersion function under a matroid constraint. In this paper, we present a PTAS for the max-sum diversification problem under a matroid constraint for distances $d(\cdot, \cdot)$ of *negative type*. Distances of negative type are, for example, metric distances stemming from the ℓ_2 and ℓ_1 norm, as well as the cosine or spherical, or Jacquard distance which are popular similarity metrics in web and image search.

Joint work with A. Cevallos and R. Zenklusen.

Non-convex compressed sensing with the sum-of-squares method

Tasuku Soma, The University of Tokyo

We consider stable signal recovery in ℓ_q quasi-norm for $0 < q \leq 1$. In this problem, given a measurement vector $y = Ax$ for some unknown signal vector $x \in \mathbb{R}^n$ and a known matrix $A \in \mathbb{R}^{m \times n}$, we want to recover $z \in \mathbb{R}^n$ with $\|x - z\|_q = O(\|x - x^*\|_q)$ from a measurement vector, where x^* is the s -sparse vector closest to x in ℓ_q quasi-norm. Although a small value of q is favorable for measuring the distance to sparse vectors, previous methods for $q < 1$ involve ℓ_q quasi-norm minimization which is computationally intractable. In this paper, we overcome this issue by using the sum-of-squares method, and give the first polynomial-time stable recovery scheme for a large class of matrices A in ℓ_q quasi-norm for any fixed constant $0 < q \leq 1$. This is a joint work with Yuichi Yoshida (National Institute of Informatics, and Preferred Infrastructure, Inc.).

The salesman's improved paths

András Sebő (CNRS, Laboratoire G-SCOP, Grenoble)

A new algorithm will be presented for the path tsp, with an improved analysis and ratio. After the starting idea of deleting some edges of Christofides' trees, we do parity correction and eventual reconnection, taking the salesman to travel through a linear program determining the conditional probabilities for some of his choices; through matroid partition of a set of different matroids for a better choice of his initial spanning trees; and through some other adventures and misadventures.

The proofs proceed by global and intuitively justified steps, where the trees do not hide the forests.

One more pleasant piece of news is that we get closer to the conjectured approximation ratio of $3/2$, and a hopefully last misadventure before finishing up this problem is that we still have to add $1/34$ to this ratio, and also for the integrality gap.

This is joint work with Anke van Zuylen.

The independent set problem on sparse graphs

Anupam Gupta (Carnegie Mellon University)

The independent set problem on graphs with maximum degree d is known to be $\Omega(d/\log^2 d)$ hard to approximate, assuming the unique games conjecture. However, the best approximation algorithm was worse by about an $\Omega(\log d)$ factor. Recently, Bansal showed how to use few rounds of the SA+ hierarchy to estimate the size of the optimal independent set to within $O(d/\log^2 d)$, essentially closing the gap. Some questions remained: could we *find* such an IS? And did we really need the SA+ lifting step?

In this talk, we will show that the standard SDP, based on the Lovasz Theta function, gives a $O(d/\log^{3/2} d)$ approximation without using any lift/project steps. I will also mention how to convert Bansal's algorithm for IS size estimation using SA+ into an approximation algorithm. Both results, just like Bansal's, are based on Ramsey theoretic results, potentially of independent interest.

This is based on joint work with Nikhil Bansal (TU Eindhoven) and Guru Guruganesh (CMU), and appeared at the STOC 2015 conference.

Robust algorithms for noisy minor-free and bounded treewidth graphs

Seeun William Umboh (TU Eindhoven)

In this talk, I will present a general approach to solve various optimization problems on noisy minor-free and bounded treewidth graphs, where some fraction of the edges have been corrupted adversarially. Our results are motivated by previous work of Magen and Moharrami, who gave a $(1 + \epsilon)$ -approximate estimation algorithm for finding a maximum independent set on noisy planar graphs, but left open the question of how to actually find one. While there are several classic approaches for planar independent set, based on recursively using separators or decomposition into k -outerplanar or bounded treewidth graphs, they break down completely with noise. Our main contribution is to design robust variants of these algorithms using LP-based techniques that also work in the noisy setting.

Joint work with Nikhil Bansal and Daniel Reichman.

Approximating directed Steiner problems via tree embedding

Bundit Laekhanukit (Weizmann Institute of Science)

In the k -edge connected directed Steiner tree (k -DST) problem, we are given a directed graph G on n vertices with edge-costs, a root vertex r , a set of h terminals T and an integer k . The goal is to find a min-cost subgraph H of G that connects r to each terminal t by k edge-disjoint r, t -paths. This problem includes as special cases the well-known directed Steiner tree (DST) problem (the case $k = 1$) and the group Steiner tree (GST) problem. Despite having been studied and mentioned many times in literature, e.g., by Feldman

et al. [SODA'09, JCSS'12], by Cheriyan et al. [SODA'12, TALG'14] and by Laekhanukit [SODA'14], there was no known non-trivial approximation algorithm for k -DST for $k \geq 2$ even in the special case that an input graph is directed acyclic and has a constant number of layers. If an input graph is not acyclic, the complexity status of k -DST is not known even for a very strict special case that $k = 2$ and $|T| = 2$.

In this talk, we present an algorithm that makes a progress toward developing a non-trivial approximation algorithm for k -DST. We present an $O(Dk^{D-1} \log n)$ -approximation algorithm for k -DST on directed acyclic graphs (DAGs) with D layers, which can be extended to a special case of k -DST on “general graphs” when an instance has a D -shallow optimal solution, i.e., there exist k edge-disjoint r, t -paths, each of length at most D , for every terminal t . For the case $k = 1$ (DST), our algorithm yields an approximation ratio of $O(D \log h)$, thus implying an $O(\log^3 h)$ -approximation algorithm for DST that runs in quasi-polynomial-time (due to the height-reduction of Zelikovsky [Algorithmica'97]). Consequently, as our algorithm works for general graphs, we obtain an $O(Dk^{D-1} \log n)$ -approximation algorithm for a D -shallow instance of the k -edge-connected directed Steiner subgraph problem, where we wish to connect every pair of terminals by k -edge-disjoint paths.

The secretary problem with a choice function

Yasushi Kawase (Tokyo Institute of Technology)

In this talk, we introduce the secretary problem with a choice function. The choice function represents the preference of the decision-maker. In this problem, the decision-maker hires some applicants, and the goal is to maximize the probability of choosing the best set of applicants defined by the choice function. When the choice function is path-independent, we provide an algorithm that succeeds with probability, at least, $1/e^k$ where k is the maximum size of the choice, and prove that this is the best possible.

Local search yields a PTAS for k -means in doubling metrics

Zachary Friggstad (University of Alberta)

The most well-known and ubiquitous clustering problem used in almost every branch of science is undoubtedly k -Means Clustering: given a set of data points X in a metric space M and a parameter k , select k centres from M to minimize the sum of the squares of the distances between points in X and their nearest centre.

In the vast majority of applications, these points lie in d -dimensional Euclidean space for some $d \geq 1$. We show that the natural p -swap local search heuristic yields a $(1 + \epsilon)$ -approximation in constant-dimensional Euclidean metrics for any constant $\epsilon > 0$. Here, p is roughly $d^d * \epsilon^{-d/\epsilon}$. Prior to this work only constant-factor approximations were known or bicriteria $(1 + \epsilon)$ -approximations that use $(1 + \epsilon) * k$ centres. A QPTAS was not even known.

This is joint work with Mohammad R. Salavatipour and Mohsen Rezapour.

Blocking arborescences

Gyula Pap (Eötvös University)

For a pair of matroids, the problem of finding a minimum (cost) set of element that blocks every common basis is NP-hard, and so are many special cases, including that of blocking every perfect matching of a graph with a minimum number of edges, even in a bipartite graph. However, there are some other special cases and related problems that may be solvable in polynomial time, including that of blocking every minimum cost arborescence, or blocking every k -arborescence, with a minimum number of arcs. There are some further results with a matroid-restricted variation, and also some problems that have remained open for the time being. This talk is an overview of some work in the past few years between Attila Bernáth, Tamás Király, and myself. For this line of research, the initial impulse came from a question raised in a paper by Naoyuki Kamiyama.

Stabilizing unstable graphs through minimum modification

Karthekeyan Chandrasekaran (University of Illinois, Urbana-Champaign)

An undirected graph G is stable if the max-fractional-matching LP (with degree and non-negativity constraints) has no integrality gap. Motivated by applications in cooperative game theory, we consider the optimization problem of achieving stability by modifying the graph to the smallest possible extent. We consider two modifications: min edge-deletion and min edge-weight-addition. We show that both these problems are NP-hard and develop approximation algorithms in certain families of graphs.

Threshold influence model for allocating advertising budgets

Naonori Kakimura (The University of Tokyo)

We propose a new influence model for allocating budgets to advertising channels. In our model, we are given a bipartite graph $G = (S, T; E)$ between commercial sources S and target customers T . Each customer $t \in T$ has a threshold θ_t . When we choose a subset X of S , a customer t is influenced if the influence he receives from X exceeds his threshold. Over the threshold model, we consider the problem of maximizing the proportion of the resulting influence to the cost spent. It is useful to measure the cost-effectiveness of a marketing campaign. We design an almost linear-time approximation algorithm to maximize the cost-effectiveness. Furthermore, we design a better-approximation algorithm based on LP rounding.

Atsushi Miyachi, Yuni Iwamasa, Takuro Fukunaga, and Naonori Kakimura, Threshold influence model for allocating advertising budgets, ICML2015, 1395–1404.

A solution to the random assignment problem with a matroidal family of goods

Satoru Fujishige (Kyoto University)

The problem of allocating a set of indivisible goods to agents in a fair and efficient manner has long been investigated in the literature. Here let us suppose that only an ordinal preference over the goods for each agent is given. A seminal paper of Bogomolnaia and Moulin (2001) shows a probabilistic serial mechanism to give a solution to the assignment problem.

In this talk we consider an extension of the random assignment problem as follows. We are given a family of sets of indivisible goods and every member of the family is realizable. In particular we consider the case when the family is a base family of a matroid.

Under the agents' ordinal preferences over goods we show an extension of the probabilistic serial mechanism to give a solution, i.e., a choice of a member (a base) of the family and an assignment. We discuss some properties of the extended probabilistic serial mechanism.

(joint work with Yoshio Sano and Ping Zhan)

Multi-label classification with pairwise relations

Seffi Naor (Technion)

Motivated by applications in multi-label learning, we introduce the metric multi-labeling problem. The objective here is to classify objects by labels while optimizing a linear cost function of both assignment costs and separation costs, which are deduced from pairwise relations between objects. Each object can be classified by multiple labels, which may have either positive or negative costs, thus departing from previous works, e.g., metric labeling. The metric multi-labeling problem is NP-hard, and we tackle it by formulating an integer program capturing the deviation from a benchmark representing an “ideal” labeling. This approach, reminiscent of the notion of regret in online learning, allows us to cope with the possible negativity of the labeling costs. We develop a tight 2-approximation algorithm for metric multi-labeling by using a counterintuitive approach that distorts the optimal likelihood values computed by our linear programming relaxation.

Joint work with Shahar Chen, Dotan Di Castro, Zohar Karnin, Liane Lewin-Eytan, and Roy Schwartz.

Approximate matchings in dynamic graph streams

Sanjeev Khanna (University of Pennsylvania)

We consider the problem of approximating a maximum matching in dynamic graph streams where the input graph is revealed as a stream of edge updates that may include both edge insertions and deletions. The goal is to design a streaming algorithm that computes an approximate matching in sublinear space. Essentially, the only known technique for designing algorithms in the dynamic streaming model is linear sketching where the algorithm maintains a linear projection of the input data.

In this talk, we will present a complete resolution of the space needed to approximate maximum matching via linear sketching in the dynamic streaming model. Combined with a recent characterization of dynamic streaming algorithms, this in fact resolves the space complexity of approximating matchings by any dynamic streaming algorithm.

We will conclude by highlighting an outstanding open problem in this area.

This is based on joint work with Sepehr Assadi, Yang Li, and Grigory Yaroslavtsev.

Constant-factor approximation for ATSP with two edge weights

Jakub Tarnawski (EPFL)

In this talk, I will discuss a recent constant factor approximation algorithm for the Asymmetric Traveling Salesman Problem on shortest path metrics of directed graphs with two different edge weights. For the case of unit edge weights, the first constant factor approximation was given recently by Svensson (FOCS 2015). This was accomplished by introducing an easier problem called Local-Connectivity ATSP and showing that a good solution to this problem can be used to obtain a constant factor approximation for ATSP. In this paper, we solve Local-Connectivity ATSP for two different edge weights. The solution is based on a flow decomposition theorem for solutions of the Held-Karp relaxation, which may be of independent interest. This is joint work with Ola Svensson and Laszlo A. Vegh.

Coloring low-discrepancy hypergraphs and promise constraint satisfaction

Venkatesan Guruswami (Carnegie Mellon University)

Given a hypergraph that admits a nearly balanced 2-coloring (of minimum possible discrepancy), we show that it is NP-hard to 2-color it without monochromatic hyperedges. I'll describe this result, and also mention two directions of follow-up work it has spurred. One is to strengthen the hardness to show intractability of coloring such hypergraphs with any constant number of colors. The other is an (ongoing) systematic study of the complexity of promise constraint satisfaction problems, aimed at proving a dichotomy theorem that characterizes pairs of Boolean constraints P, Q for which Q -satisfying a P -satisfiable instance is tractable (the case $P = \text{Majority}$ and $Q = \text{Or}$ is one such example, which states that k -SAT is easy when there's an assignment satisfying a majority of literals in each clause).

[Touches on several works, joint with Per Austrin, Johan Hastad, Euiwoong Lee, and Joshua Brakensiek.]

An improve approximation for weighted completion time on unrelated machines

Nikhil Bansal (TU Eindhoven)

I will talk about a recent $(3/2 - \epsilon)$ -approximation for the problem of schedul-

ing jobs on unrelated machines so as to minimize the sum of weighted completion times. This improves upon the long-standing bound of $3/2$. The result is based on a new lift-and-project based SDP relaxation for the problem and a new general bipartite-rounding procedure that produces an assignment with certain strong negative correlation properties.

Joint work with Aravind Srinivasan and Ola Svensson

Combinatorial algorithms for some multifold problems and related network designs

Hiroshi Hirai (The University of Tokyo)

We address the following two multifold problems: (i) The minimum cost node-demand multifold problem. (ii) The maximum node-capacitated free multifold problem. For (i), we develop an $O(n \log(nAC)MF(kn, km))$ time algorithm for a network of n nodes, m edges, k terminals, the maximum A of edge-cost, and the sum C of edge-capacity, where $MF(n', m')$ is the time complexity of solving the maximum flow problem for a network of n' nodes and m' edges. For (ii), we develop an $O((m \log k)MSF(n, m, 1))$ time algorithm for a network of n nodes, m edges, and k terminals, where $MSF(n, m, h)$ is the time complexity of solving the maximum submodular flow problem on a network of n nodes, m edges, and time complexity h of computing the exchange capacity of the defining submodular function. By Fujishige-Zhang algorithm for submodular flow, we obtain an $O(mn^3 \log k)$ time algorithm for (ii). Our algorithms are designed on the basis of a developing theory of discrete convex functions on certain graph structures, and bring “ellipsoid-free” combinatorial implementations to approximation algorithms of related network design problems.

Approximation algorithms for inventory problems with submodular or routing costs

Viswanath Nagarajan (University of Michigan)

We consider inventory optimization problems over a horizon of T periods that involve two kinds of costs (1) an ordering cost to replenish items and (2) a holding cost to store items. When the ordering cost satisfies a certain natural property and the holding cost is a polynomial function, we obtain an $O(\log T / \log \log T)$ approximation algorithm. This implies the first sublogarithmic approximation ratio for the inventory-routing and submodular-joint-replenishment problems, under linear holding costs.

Recent developments in robust network flows

Jannik Matuschke (TU Berlin)

Robust network flows are a concept for dealing with uncertainty and failures in the network infrastructure. One of the most basic models is the Maximum Robust Flow problem: Given a network and an integer k , the task is to find a path flow maximizing the guaranteed value of surviving flow after failure of any k arcs in the network. The complexity of this problem appeared to have

been settled almost a decade ago: Aneja et al. (2001) showed that it can be solved efficiently when $k = 1$, while a paper by Du and Chandrasekaran (2007) claimed NP-hardness for any constant $k > 1$. We point to a fatal flaw in this latter result, and give a new strong NP-hardness proof for the case that k is unbounded. This leaves the complexity of the problem for constant number of arc failures open once again. We will also discuss a model where flow can be re-routed after failure of links and provide hardness results and algorithms for this setting.

This is joint work with Yann Disser, Tom McCormick, and Gianpaolo Oriolo.

4 List of Participants

- Hyung-Chan An (Yonsei University)
- Satoru Iwata (University of Tokyo)
- Zachary Friggstad (University of Alberta)
- Sanjeev Khanna (University of Pennsylvania)
- Alina Ene (University of Warwick)
- Seeun William Umboh (TU Eindhoven)
- Anupam Gupta (Carnegie Mellon University)
- Neil Olver (VU Amsterdam)
- András Sebő (CNRS, Laboratoire G-SCOP, Grenoble)
- Hiroshi Hirai (University of Tokyo)
- Kenjiro Takazawa (Kyoto University)
- Alantha Newman (CNRS, Grenoble)
- Viswanath Nagarajan (University of Michigan)
- Tobias Mömke (Saarland University)
- Bundit Laekhanukit (Weizmann Institute of Science)
- Debmalya Panigrahi (Duke University)
- Seffi Naor (Technion)
- Nikhil Bansal (TU Eindhoven)
- Friedrich Eisenbrand (EPFL)
- Tamás Király (ELTE Budapest)
- Gyula Pap (Eötvös University)
- Venkatesan Guruswami (Carnegie Mellon University)
- Karthekeyan Chandrasekaran (University of Illinois, Urbana-Champaign)

- Jannik Matuschke (TU Berlin)
- Naonori Kakimura (University of Tokyo)
- Kristóf Bérczi (MTA-ELTE, Budapest)
- Jakub Tarnawski (EPFL)
- Satoru Fujishige (Kyoto University)
- Tasuku Soma (University of Tokyo)
- Yasuhi Kawase (Tokyo Institute of Technology)
- Yutaro Yamaguchi (Osaka University)
- Ken-ichi Kawarabayashi (National Institute of Informatics)
- Takuro Fukunaga (National Institute of Informatics)
- R. Ravi (Carnegie Mellon University)
- László Végh (London School of Economics)