ISSN 2186-7437

NII Shonan Meeting Report

No. 2011-1

Graph Algorithm and Combinatorial Optimization

Satoru Iwata
Ken-ichi Kawarabayashi

February 13–18, 2011

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

Graph Algorithm and Combinatorial Optimization

Organizers:
Satoru Iwata (RIMS, Kyoto University)
Ken-ichi Kawarabayashi (NII)

February 13–18, 2011

Structures that can be represented as graphs are ubiquitous, and many practical problems can be represented by graphs. The link structure of a web in Internet could be represented by a directed graph: the vertices are the web pages available at the website and a directed edge from page A to page B exists if and only if A contains a link to B. A graph structure can be extended by assigning a weight to each edge of the graph.

Networks have also many uses in the practical side of graph theory, network analysis (for example, to model and analyze traffic networks). Within network analysis, the definition of the term "network" varies, and may often refer to a simple graph.

A similar approach can be taken to problems in travel, biology, computer chip design, and many other fields. Development of algorithms to handle graphs is therefore of major interest in computer science.

These applications of graphs often gives rise to optimization. Basic optimization problems on graphs, including the shortest path, maximum flow, minimum spanning tree problems allow efficient exact algorithms. The algorithmic developments of these problems have led to the theory of combinatorial optimization, combined with polyhedral combinatorics, matroids and submodular functions.

On the other hand, most practical optimization problem on graphs such as the traveling salesman, stable set, maximum cut problems are NP-hard. Approximation algorithms for these NP-hard combinatorial optimization problems have been investigated extensively for a couple of decades. Design of approximation algorithms often requires deep insights from structural graph theory and polyhedral combinatorics.

The purpose of this workshop is to bring experts in graph algorithm and combinatorial optimization to share ideas, and to stimulate joint projects.

Overview of Talks

Properties of sparse graphs

Zdeněk Dvořák, Charles University, Prague

My research concerns mainly properties of sparse graphs, especially graphs close to planar and more generally, classes of graphs with bounded expansion and nowhere-dense graph classes (which are the natural definitions of "sparse graphs" for properties definable in first-order logic, including many important graph classes – e.g., proper minor-closed graph classes, graphs with bounded maximum degree, or graphs drawn in a fixed surface with a bounded number of crossings on each edge). Together with D. Král' and R. Thomas, we gave a linear-time algorithm for deciding first-order properties for graphs with bounded expansion, generalizing and improving all previously known algorithms for this problem on sparse graphs.

Another interesting problem in this area that I would like to consider is the following.

Problem 1 *Which graph classes $\mathcal{C}$ have the following property: there exists a function $f : Z^+ \to Z^+$ such that for every $G \in \mathcal{C}$ and for every integer $p > 1$, G has a (not necessarily proper) coloring by p colors such that the union of every $p-1$ color classes has treewidth at most $f(p)$?*

Every class of graphs with this property admits approximation and FPT algorithms for variety of important problems. It is known that proper minor-closed graph classes have this property; how much can this claim be generalized?

I also work on various graph coloring problems; recently, together with B. Lidický, B. Mohar and R. Škrekovski, we considered the choosability of graphs drawn in plane with crossings (or other obstructions, like precolored vertices) far apart.

On stable matchings and flows

Tamás Fleiner, Budapest University of Technology and Economics

We describe a flow model that is related to ordinary network flows the very same way as stable matchings are related to bipartite matchings. That is, each vertex of the network represents an agents trading with some "stuff" and having preferences on the trading partners. A stable trading scheme is a flow and stability can be defined in such a way that no agent can improve her situation by moving some trade from a less preferred partner to a more preferred one.

One can prove that there always exists a stable flow and it is also possible to generalize the lattice structure of stable marriages to stable flows. The main tool is a reduction of the stable flow problem to stable allocations, that is, to stable b-matchings.

However, there is a a little problem with the definition of stable flows. It might happen that a stable flow allows an unsaturated cycle in the underlying network, and in a sense this is a reason of unstability: if every agent sells and buys more along the cycle then everyone is happier. So a flow is *completely stable* if it is stable and no such cycle exists. One can easily find a network with preferences in such a way that no completely stable flow exists. It is an interesting problem to decide if there is an efficient algorithm that finds a completely stable flow if such exists.

Problems in Submodular Function Minimization That I Still Want to Solve

Satoru Fujishige, RIMS, Kyoto University

I listed some problems in submodular function minimization (SFM) in my paper [7], from which I excerpt (Item 4 is added):

1. Another Framework for SFM: Inner expressions for base polyhedra
   The common feature of Schrijver's [12] and the IFF algorithm [11] is the use of convex combination of extreme bases. This is the framework of Cunningham's approach [2, 3]. Can we devise an algorithm without convex combinations?

   Given some extreme bases y_i $(i \in I)$, if they are affinely independent, the set of points that can be expressed as a convex combination of the extreme bases forms a simplex. We can also get a polyhedron by translating each facet inequality of the base polyhedron so that it is valid and tight for some given extreme base y_i. In other words and more precisely, define a set function $h : 2^V \to \mathbf{R}$ by

   $$h(X) = \max\{y_i(X) \mid i \in I\} \quad (X \subseteq V), \tag{1}$$

   and a polyhedron $\mathrm{Com}(P)$ with $P = \{y_i \mid i \in I\}$ by

   $$\mathrm{Com}(P) = \{x \in \mathbf{R}^V \mid \forall X \subseteq V : \ x(X) \le h(X),\ x(V) = h(V)\}, \tag{2}$$

   where note that $h(V) = f(V)$. We call $\mathrm{Com}(P)$ the *combinatorial hull* of P ([6]). We can easily see that by definition the combinatorial hull is contained in the base polyhedron and contains the convex hull of P. The function h that expresses the combinatorial hull is a subadditive set function smaller than f. It should, however, be noted that the combinatorial hull is not easy to handle. Even the membership in a combinatorial hull of two bases in general is hard to test. To overcome this difficulty we can consider a slightly weaker approach along the idea of combinatorial hull as follows. For a pair of bases y_1 and y_2, if for some $u \in V$ we have

   $$y_1(u) < y_2(u), \qquad y_1(v) \ge y_2(v) \quad (v \in V \setminus \{u\}), \tag{3}$$

   then the combinatorial hull of the two is given by

   $$\mathrm{Com}(\{y_1, y_2\}) = \{z \in \mathbf{R}^V \mid z(V) = f(V),\ y_1(u) \le z(u) \le y_2(u),\ y_1(v) \ge z(v) \ge y_2(v)\ (v \in V \setminus \{u\})\}. \tag{4}$$

   Hence we need only comparisons and additions to test membership in the combinatorial hull of such two bases. We might be able to recursively apply this technique to show a membership in the base polyhedron by using a set of extreme bases and by repeatedly making combinatorial hulls of two bases, possibly non-extreme bases generated before. Combinatorial hull employed in such a way could be a basic tool to obtain another framework for (fully) combinatorial SFM algorithms.

2. Another Framework for SFM: Outer expressions for base polyhedra
Let us now consider outer expressions for a base polyhedron. Each facet inequality is a linear approximation of the base polyhedron. Also, a tangent cone at an extreme base can be regarded as a quadratic approximation. Testing membership in such a tangent cone can be solved by finding a maximum flow or a minimum cut in a network with its underlying graph being the poset associated with the extreme base (see [1, 5]). There may be some successive QP-type algorithm for SFM by means of quadratic approximation.

3. A Descent Method
I draw your attention to a descent method for SFM due to Fujishige and Iwata [10]. Note that the membership problem for base polyhedra is equivalently given as follows: Discern whether f is nonnegative and, if f is not nonnegative, give a set X such that $f(X) < 0$, where f can be any submodular function. We can solve the SFM problem by $\mathrm{O}(n^2)$ descent steps by using an oracle for the membership problem. This algorithm is fully combinatorial modulo the membership oracle. Can we make it fully combinatorial and polynomial without the membership oracle? Or can we devise a fully combinatorial polynomial algorithm for testing membership in base polyhedra that runs faster than SFM algorithms? The technique of combinatorial hull mentioned above could be a good candidate for resolving these questions.

4. Practically Efficient SFM Algorithm and Its Complexity
A practically efficient SFM algorithm is given by the use of Wolfe's algorithm [13] to compute the minimum-norm base (see [4, 8, 9]). It runs very fast [9] but its complexity is still unknown.

References

[1] R. E. Bixby, W. H. Cunningham, and D. M. Topkis (1985): Partial order of a polymatroid extreme point, *Mathematics of Operations Research*, **10**, 367–378.

[2] W. H. Cunningham (1984): Testing membership in matroid polyhedra, *Journal of Combinatorial Theory*, **B36**, 161–188.

[3] W. H. Cunningham (1985): On submodular function minimization, *Combinatorica*, **5**, 185–192.

[4] S. Fujishige (1984): Submodular systems and related topics, *Mathematical Programming Study*, **22**, 113–131.

[5] S. Fujishige (1984): On the subdifferential of a submodular function. *Mathematical Programming*, **29**, 348–360.

[6] S. Fujishige (2000): Combinatorial hull and submodular function minimization, 17th International Symposium on Mathematical Programming, Atlanta, GA USA, August 7-11, 2000.

[7] S. Fujishige (2003): Submodular function minimization and related topics. *Optimization Methods and Software*, **18**, 167–180.

[8] S. Fujishige (2005): *Submodular Functions and Optimization* Second Edition, *Annals of Discrete Mathematics*, **58**, Elsevier.

[9] S. Fujishige and S. Isotani (2011): A submodular function minimization algorithm based on the minimum-norm base. *Pacific Journal of Optimization*, **7**, 3–17.

[10] S. Fujishige and S. Iwata (2002): A descent method for submodular function minimization. *Mathematical Programming*, Ser. A **92**, 387–390.

[11] S. Iwata, L. Fleischer, and S. Fujishige (2001): A combinatorial strongly polynomial algorithm for minimizing submodular functions, *Journal of ACM*, **48**, 761–777.

[12] A. Schrijver (2000): A combinatorial algorithm minimizing submodular functions in strongly polynomial time. *Journal of Combinatorial Theory*, Ser. B **80**, 346–355.

[13] P. Wolfe (1976): Finding the nearest point in a polytope, *Mathematical Programming*, **11**, 128–149.

Submodular system partition problem

Takuro Fukunaga, Kyoto University

The submodular system k-partition problem is a problem of partitioning a given finite set V into k non-empty subsets $V_1, V_2, \ldots, V_k$ so that $\sum_{i=1}^{k} f(V_i)$ is minimized where f is a non-negative submodular function on V, and k is an integer at least two. This problem contains the graph and hypergraph k-cut problems as special cases. In particular, the graph k-cut problem can be formulated by symmetric submodular functions.

I am interested in computational complexity of the hypergraph k-cut problem and the submodular system k-partition problem with symmetric/non-symmetric sumodular functions. The graph k-cut problem is known to be NP-hard if k is a part of inputs, and hence all the above problems are NP-hard in that case. However the graph k-cut problem admits polynomial-time exact algorithms if k is a fixed constant. My question is whether there exist polynomial-time exact algorithms for the hypergraph k-cut problem and the submodular system k-partition problem when k is a fixed constant. So far, we have obtained the following results:

1. Submodular system k-partition problem with general submodular functions: Okumoto, Fukunaga & Nagamochi (2009) presented a divide-and-conquer algorithm. The algorithm is an exact algorithm for $k = 3$, a $(k-1)/2$-approximation algorithm for $4 \leq k \leq 6$, and a $(k+1-2\sqrt{k-1})$-approximation algorithm for $k \geq 7$.

2. Submodular system k-partition problem with symmetric submodular functions: It was shown by Zhao, Nagamochi & Ibaraki (2001), and Queyranne (1999) independently that a greedy algorithm achieves 2-approximation. I have observed that the algorithm due to Okumoto, Fukunaga & Nagamochi (2009) is an exact algorithm for $k = 4$ in this case.

3. Hypergraph k-cut problem: For $k = 3$, a polynomial-time exact algorithm was proposed by Xiao (2008). Fukunaga (2010) showed that if both k and the maximum size of hyperedges are fixed, it can be solved in polynomial-time. Okumoto, Fukunaga & Nagamochi (2009) observed that this problem is contained by the terminal k-vertex cut in graphs, for which there exists a $(2 - 2/k)$-approximation algorithm.

Prize-collecting Clustering and Algorithmic Applications

Mohammad Hajiaghayi

Grouping a set of items into clusters according to their similarities is called-clustering. It is now a common technique in widely different applications in, e.g., bioinformatics, data mining, machine learning, and social network analysis. Considerable effort has been put into study of clustering techniques in recent years.

In this thesis, we introduce a new clustering paradigm, in which items are vertices of a graph. Vertices have their own budgets and the goal is to cluster them such that the cost of (connections in) each cluster can be payed by the budget of its participants. Furthermore, we want vertices in different clusters be in some sense far from each other.

We propose and analyze algorithms for two similar problems in this paradigm. The resulting guarantees lead to resolution of several network design questions. We improve the approximation ratio for prize-collecting Steiner tree and TSP. In addition, we present polynomial-time approximation schemes (PTAS's) for planar Steiner forest, planar multiway cut, and planar prize-collecting Steiner tree.

Tree metrics and edge-disjoint S-paths

Hiroshi Hirai, University of Tokyo

Given an undirected graph $G = (V, E)$ with a terminal set S, a terminal weight $\mu : \binom{S}{2} \to \mathbf{Z}_+$, and an edge-cost $a : E \to \mathbf{Z}_+$, the μ-weighted minimum-cost edge-disjoint S-paths problem (μ-CEDP) is to maximize $\sum_{P\in\mathcal{P}} \mu(s_P, t_P) - a(P)$ over all edge-disjoint sets $\mathcal{P}$ of S-paths, where s_P, t_P denote the ends of P and $a(P)$ is the sum of edge-cost $a(e)$ over edges e in P.

Our main result is a complete characterization of terminal weights μ for which μ-CEDP is tractable and admits a combinatorial min-max theorem for every graph. We prove that if μ is a tree metric, then μ-CEDP is solvable in polynomial time and has a combinatorial min-max formula, which extends Mader's edge-disjoint S-paths theorem and its minimum-cost version by Karzanov. Our min-max theorem solves the dual half-integrality conjecture in the minimum-cost edge-disjoint S-paths raised by Karzanov as a special case. We also prove that the cost-less version (μ-EDP) is NP-hard if μ is not a truncated tree metric, where a truncated tree metric is a weight function represented as pairwise distances among balls in a tree. On the other hand, μ-EDP for a truncated tree metric μ reduces to μ'-CEDP for a tree metric μ'. Thus our result is best possible unless $P = NP$. As an application, we get a good approximation algorithm for μ-EDP with "near" tree metric μ by utilizing results from the theory of low-distortion embedding. This is a joint work with Gyula Pap (Eötvös University, Budapest).

A Tighter Insertion-based Approximation of the Crossing Number

Petr Hliněný, FI MU Brno, CZ

The *crossing number* problem of a graph G is to find a drawing of G (in the plane) minimizing the number of pairwise edge crossings in it. The *multiple edge insertion* problem (MEI) of a graph G, and edge set F not in G, is to find a crossing-minimizing drawing of $G+F$ such that the subdrawing restricted to G is planar. We provide a polynomial time algorithm that approximates the solution to MEI up to an additive factor (depending on $\Delta(G)$ and $|F|$ only), which in turn gives an approximation of the crossing number of $G+F$ up to a multiplicative factor. Our algorithm is simple and practically implementable.

The crossing number problem is NP-complete, even in the case of almost planar graphs ($G+e$ where G is planar) by Cabello–Mohar [SoCG 2010]. The single edge insertion problem has got an exact linear-time algorithm by Gutwenger et al. [Algorithmica 2005], while the general MEI with arbitrary F is NP-complete by Ziegler. A special case of MEI, namely the single vertex insertion problem, has also a polynomial-time exact solution by Chimani et al [SODA 2009]. The close algorithmic connection between the aforementioned two problems was established by Hliněný–Salazar [GD2006] and Cabello–Mohar [GD 2008] with proving that a solution to edge insertion approximates the crossing number of $G+e$.

Chimani et al [GD 2008] then proved that vertex insertion approximates the crossing number of apex graphs, and that also a solution to MEI generally approximates the crossing number of $G+F$. The drawback of the latter is that no exact efficient solution to MEI is known even with fixed $|F|>1$. Chuzhoy et al [SODA 2011] have recently provided a combined approximation algorithm for both the MEI and crossing problems on planar G and F. Our new algorithm provides better approximation guarantees for both, and is simpler at the same time.

This is a joint work with Markus Chimani of Friedrich-Schiller-University Jena, Germany.

Finding 2-factors covering 3- and 4-edge cuts in bridgeless cubic graphs

Satoru Iwata, Research Institute for Mathematical Sciences, Kyoto University

A famous theorem of Petersen states that every bridgeless cubic graph contains a perfect matching, and hence a 2-factor. In fact, such a graph has a 2-factor that covers all the 3-edge cuts. A recent paper of Kaiser and Škrekovski (2008) shows that every bridgeless cubic graph has a 2-factor that covers all the 3- and 4-edge cuts. In this work, we provide an efficient algorithm to find such a 2-factor. Using this algorithm as a preprocess, we also devise a simple

6/5-approximation algorithm for finding a minimum 2-edge-connected spanning subgraph in 3-edge-connected cubic graphs. This is a joint work with Sylvia Boyd and Kenjiro Takazawa.

Highly connected rigidity matroids have unique underlying graphs

Tibor Jordán, Department of Operations Research, Eötvös University, Budapest, Hungary

Let $\mathcal{M}$ be a d-dimensional generic rigidity matroid of some graph G. We consider the following problem, posed by Brigitte and Herman Servatius: is there a (smallest) integer k_d such that the underlying graph G of $\mathcal{M}$ is uniquely determined, provided that $\mathcal{M}$ is k_d-connected? Since the one-dimensional generic rigidity matroid of G is isomorphic to its cycle matroid, a celebrated result of Hassler Whitney implies that $k_1 = 3$. We extend this result by proving that $k_2 \leq 11$. To show this we prove that (i) if G is 7-vertex-connected then it is uniquely determined by its two-dimensional rigidity matroid, and (ii) if a two-dimensional rigidity matroid is $(2k-3)$-connected then its underlying graph is k-vertex-connected.

We also prove the reverse implication: if G is a k-connected graph for some $k \geq 6$ then its two-dimensional rigidity matroid is $(k-2)$-connected. Furthermore, we determine the connectivity of the d-dimensional rigidity matroid of the complete graph K_n, for all pairs of positive integers d, n. (Joint work with Viktória Kaszanitzky.)

Robust Independence Systems

Naonori Kakimura, University of Tokyo

An independence system is one of the most fundamental combinatorial concepts, which includes a variety of objects in graphs and hypergraphs such as matchings, stable sets, and matroids. We discuss the robustness for independence systems, which is a natural generalization of the greedy property of matroids. For a real number $\alpha > 0$, a set $X \in \mathcal{F}$ is said to be *α-robust* if for any k, it includes an α-approximation of the maximum k-independent set, where a set Y in $\mathcal{F}$ is called k-independent if the size $|Y|$ is at most k. In this talk, we show that every independence system has a $1/\sqrt{\mu(\mathcal{F})}$-robust independent set, where $\mu(\mathcal{F})$ denotes the *exchangeability* of $\mathcal{F}$. Our result contains a classical result for matroids and the ones of Hassin and Rubinstein [SIAM Disc. Math. '02] for matchings and Fujita, Kobayashi, and Makino [ESA '10] for matroid 2-intersections, and provides better bounds for the robustness for many independence systems such as b-matchings, hypergraph matchings, matroid p-intersections, and unions of vertex disjoint paths. Furthermore, we provide bounds of the robustness for nonlinear weight functions such as submodular and convex quadratic functions. We also extend our results to independence systems in the integral lattice with separable concave weight functions. This is a joint work with Kazuhisa Makino.

Stable Matching Models with Lower Quotas and Discrete Convex Analysis

Naoyuki Kamiyama, Chuo University

In this talk, we consider a relation between the following two recent progresses on stable matching models: the extension to the framework of discrete convex analysis and the introduction of lower quotas. A discrete-concave stable matching model presented by Eguchi, Fujishige and Tamura '03 has great generality, and it includes many variants of the classical hospital/resident problem. However, this model does not include stable matching models with lower quotas which are presented by Huang '09 and Kamiyama '10. So the open problem can be described as follows.

Problem 1 *Can we extend the known results for stable matching models with lower quotas to the framework of discrete convex analysis?*

Minimum k-way cut of bounded size is fixed-parameter tractable

Ken-ichi Kawarabayashi, National Institute of Informatics, Tokyo, Japan

We consider a the minimum k-way cut problem for unweighted graphs with a size bound s on the number of cut edges allowed. Thus we seek to remove as few edges as possible so as to split a graph into k components, or report that this requires cutting more than s edges. We show that this problem is fixed-parameter tractable (FPT) in s. More precisely, for $s = O(1)$, our algorithm runs in quadratic time while we have a different linear time algorithm for planar graphs and bounded genus graphs.

Our tractability result stands in contrast to known W[1] hardness of related problems. Without the size bound, Downey et al. [2003] proved that the minimum k-way cut problem is W[1] hard in k even for simple unweighted graphs. Downey et al. asked about the status for planar graphs. Our result implies tractability in k for the planar graphs since the minimum k-way cut of a planar graph is of size at most $6k$ (in fact, the size is $f(k)$ for any bounded average degree graphs for some fixed function f of k. This class includes bounded genus graphs, and simple graphs with an excluded minor).

A simple reduction shows that vertex cuts are at least as hard as edge cuts, so the minimum k-way vertex cut is also W[1] hard in terms of k. Marx [2004] proved that finding a minimum k-way vertex cut of size s is also W[1] hard in s. Marx asked about the FPT status with edge cuts, which we prove tractable here. We are not aware of any other cut problem where the vertex version is W[1] hard but the edge version is FPT.

Joint work with Mikkel Thorup

Multi-Layered Video Streaming with Network Coding

Zoltán Király, Eötvös University, Budapest

In multi-layered video streaming receivers may have different quality requirements. In multi-resolution coding a layer is **valuable** for a terminal node, only if this node receives all the layers with higher importance. Let $D = (V, A)$ be a directed acyclic graph with a single source node s and with unit capacity arcs. Actual packets of the k layers correspond to a set of messages $\mathbf{M} = (M_1, M_2, \ldots M_k)$ represented by members of a finite field $\mathbb{F}_q$ of size q. The task is to multicast $\mathbf{M}$ from s. The idea of network coding is to transmit linear combinations of messages in $\mathbf{M}$ on the arcs. Such a linear combination on arc uv can be represented by the vector of the coefficients: $\mathbf{c}(uv) \in \mathbb{F}_q^k$, where $\mathbf{c}(uv)$ is a linear combination of $\{\mathbf{c}(wu) | \ wu \in A\}$. A **request** of a node can be the first i layers, where $0 \leq i \leq k$. A node t with request i receives (can decode) its request, if the first i unit vectors are in the span of $\{\mathbf{c}(ut) | \ ut \in A\}$. A **demand** τ is an ordered set of disjoint subsets of $V \setminus s$ denoted by $\tau = (T_1, T_2, \ldots, T_k)$, where T_i denotes the set of nodes with request i. $T = T_1 \cup \ldots \cup T_k$ is the set of terminals.

We show that the the following two problems are NP-complete.

a) Given D, a demand $\tau = (T_1, T_2)$ and an integer K, decide whether there exists a network code, where at least K terminals can decode their requests.

b) Given D and a demand $\tau = (T_1, \emptyset, T_3)$, decide whether there exists a network code, where all terminals can decode their requests.

On the other hand, for two layers ($k = 2$), we give an algorithm that constructs a network code, where every terminal receives the first layer, and under this assumption, the unique maximum cardinality subset of T_2 receives the second layer.

For three layers we give a heuristic, where every terminal receives the first layer, the most possible terminals in $T_2 \cup T_3$ receive the second layer, and some terminals in T_3 receive the third layer. For this heuristic we developed a distributed algorithm that decides for every terminal in $T_2 \cup T_3$, whether it can decode the second layer under the assumption, that every terminal receives the first layer.

This is joint work with Erika Kovács.

A polynomial-time approximation scheme for planar multiterminal cut

Philip Klein, Brown University

The *multiterminal cut* problem is as follows: given an undirected graph with edge-costs and a subset of nodes (called the *terminals*), find a minimum-cost subset of edges whose removal disconnects each terminal from the others. It generalizes the min *st*-cut problem but is NP-hard for planar graphs and APX-hard for general graphs. We give a polynomial-time approximation scheme for this problem on planar graphs.

This is joint work with MohammadHossein Bateni, MohammadTaghi Hajiaghayi, and Claire Mathieu.

Breaking $O(n^{1/2})$-approximation algorithms for the edge-disjoint paths problem with congestion two

Yusuke Kobayashi, University of Tokyo

In the maximum edge-disjoint paths problem, we are given a graph and a collection of pairs of vertices, and the objective is to find the maximum number of pairs that can be routed by edge-disjoint paths. A c-approximation algorithm for this problem is a polynomial time algorithm that finds at least OPT/c edge-disjoint paths, where OPT is the maximum possible. Currently, an $O(n^{\frac{1}{2}})$-approximation algorithm is best known for this problem even if a congestion of two is allowed, i.e., each edge is allowed to be used in at most two of the paths.

In this paper, we give the first result that breaks the $O(n^{\frac{1}{2}})$-approximation with congestion two. Specifically, we give a randomized $O(n^{\frac{3}{7}})$-approximation algorithm. Our framework for this algorithm is more general in a sense. Indeed, we have two ingredients which also work for the edge-disjoint paths problem (with congestion one) if the following conjecture is true.

Conjecture: There is a (randomized) polynomial-time algorithm for finding $\Omega(\mathrm{OPT}^{\frac{1}{p}}/\beta(n))$ edge-disjoint paths connecting given terminal pairs, where β is a polylogarithmic function.

Having made this conjecture, we prove the following.

1. Assuming the above conjecture for some $p > 1$, for some absolute constant $\alpha > 0$, we show that by using Rao-Zhou's algorithm, we can give a randomized $O(n^{\frac{1}{2}-\alpha})$-approximation algorithm for the edge-disjoint paths problem (with congestion one).
2. Based on the Räcke decomposition and Chekuri-Khanna-Shepherd well-linked set, we show that there is a randomized algorithm for finding $\Omega(\mathrm{OPT}^{\frac{1}{4}})$ edge-disjoint paths connecting given terminal pairs with congestion two (thus confirming the above conjecture if we allow congestion to be two).

All of our results still hold for the vertex-disjoint paths problem as well, i.e., paths are not edge-disjoint, but vertex-disjoint case.

This is joint work with Ken-ichi Kawarabayashi.

Exponentially many perfect matchings in cubic bridgeless graphs

Daniel Král', Charles University

We show that every cubic bridgeless graph with n vertices has at least $2^{n/3656}$ perfect matchings. This confirms a conjecture of Lovász and Plummer.

The result is joint with L. Esperet, F. Kardos, A. King and S. Norine.

Computing graph connectivity by network coding

Lap Chi Lau, The Chinese University of Hong Kong

We present algebraic algorithms for computing edge-connectivities in directed graphs. Using ideas from network coding, we reduce the problem to solving systems of linear equations and computing the rank of the solution. This allows us to use tools from linear algebra to obtain faster algorithms to compute single-source edge-connectivities and all-pairs edge-connectivites.

1. For directed acyclic graphs, there is an $O(mn + mn^{\omega-1})$ algorithm to compute the edge-connectivities from a source to all other vertices, where $\omega \approx 2.376$ is the matrix multiplication exponent. Interestingly, superconcentrators are used in order to solve the linear equations faster.

2. For bounded degree planar graphs, there is an $O(n^{\omega/2})$ algorithm to compute the edge-connectivities from a source to all other vertices. This is based on the recent result of Alon and Yuster on "Solving linear equations through nested dissection".

3. For general directed graphs, there is an $O(m^{\omega})$ algorithm to compute all-pairs edge-connectivities. This is faster than the best known combinatorial algorithm.

4. For graphs with "good" separators (e.g. bounded degree fixed minor free graphs), there is an $O(n^{1+\omega/2})$ algorithm to compute all-pairs edge-connectivities. The algorithm is based on a faster algorithm to compute the inverse of a well-separable matrix.

The use of superconcentrator can also be applied to obtain faster algorithm for edge splitting-off in directed graphs. This is joint work with Ho Yee Cheung and Kai Man Leung.

Combinatorial Algorithms for TDI Systems

Tom McCormick, Sauder School of Business, UBC

There are many results showing that linear programs with specially structured right-hand sides (often some form of submodularity) are *totally dual integral (TDI)*, and so have guaranteed integral optimal solutions (with integral data). However, corresponding polynomial-time combinatorial algorithms for such problems are often lacking. Together with Maren Martens, I showed such an algorithm for Hoffman's *Weighted Abstract Flow (WAF)* model at IPCO 2008, and with Britta Peis I showed such an algorithm for a version of Hoffman's Lattice Polyhedron model that we call *Weighted Abstract Cut Packing* that will appear at IPCO 2011. Both algorithms are based on the Primal-Dual Algorithm framework and use capacity scaling to achieve weak polynomiality. Together with Fujishige, I also found such an algorithm for minimizing bisubmodular functions (BSFM).

It is natural to wonder if these algorithms can be further extended: can the same methods solve over general lattice polyhedra with 0, ± 1 matrices? What about WAF with 0, ± 1 matrices? What about Schrijver's general framework for TDI problems?

An interesting open problem is to find a polynomial combinatorial algorithm for optimizing over the *subtour elimination polytope* associated with the Traveling Salesman Problem (TSP), which was named as one of 10 important unsolved problems in TSP by Bill Cook. There is some indication that this problem could be solved by adapting some of the same techniques that were used by IFF to solve submodular function minimization (SFM).

My main research interests

Gianpaolo Oriolo, Università Tor Vergata, Roma

My main interest are in the area of packing and network design.

Packing problems. I am particularly interested in stable (or independent) set problems, both combinatorial and polyhedral aspects. The problem that puzzles me most is the maximum weighted stable set on claw-free graphs. Claw-free graphs are a superclass of line graphs, therefore the maximum weighted stable set on claw-free graphs is a (fundamental) generalization problem of the matching problem. In a joint work with Yuri Faenza and Gautier Stauffer we recently managed to design an $O(n^3)$ algorithm for the solution of this problem (Gautier submitted this talk to our workshop). Note that, following matching algorithms, the maximum weighted stable set on line graphs can be solved in $O(n^2 \log n)$, so the complexity of our algorithm for claw-free graphs *almost* matches that complexity. One question that is open and I find interesting is whether we can "close" this gap, which I think should be possible at least for the class of *quasi-line* graphs: a graph is quasi-line if the neighborhood of any vertex can ve partitioned into two cliques. Note that quasi-line graphs are a superclass of line graphs and a subclass of claw-free graphs.

Other open questions that I find attracting concern the stable set polytope of claw-free graphs. This polytope generalizes of the matching polytope, but it seems quite more involved. We have a complete and reasonably nice description for it for the class of quasi-line graphs (this is the solution to the so-called Ben Rebea's Conjecture), but we don't have a *minimal* description. A first open question is whether we can we get such a description, and more in general I wonder which combinatorial structures are "responsible" for the non-matching like facets of this polytope. Another open question concerns the description of this polytope for the class of claw-free graphs. There are some results in this direction, but either they are very technical or they are based on extended formulations. Still a very nice fact about the stable set polytope of claw-free graphs has been observed by Calvillo. The question is then the following: can be build a reasonably nice description of this polytope upon Calvillo's Theorem?

Network design. I am interested in the design of approximation algorithms for the design (i.e., min cost capacity installation) of *robust* and *resilient* networks. To me, a resilient network is a network that is still operational even if some failure event happen; while a robust network is a network that is operational under different patterns of traffic demands. In particular, I've been recently working on the VPN problem, that is the problem of installing min-cost capacity on an undirected network, as to satisfy a suitable (polyhedral) set of non-simulatenous traffic demands under some routing constraints. This problem attracted quite a few researchers in the last years, especially from the Computer Science community. In a joint work with Nicola Apollonio, Fabrizio Grandoni and András Sebo, we recently we initiated a polyhedral study of this problem that led to simpler proofs of some known results as well as to deeper understanding and some sharpening of these results (András submitted this talk to our workshop).

Finally let me mention a related *simple* open problem that can be attacked without being a network design expert. Suppose that we are given an undirected network, a source r and set of terminals $W_1, W_2, \ldots W_k$. Each set of terminals W_i determines the following set of traffic demands D_i: r has to send one unit of flow to each terminal in W_i. The question: is there a (constant factor) approximation algorithm for the problem of finding a min-cost integral capacity installation as to route (non-simultaneously) each set of traffic demands $D_1, \dots D_k$?

Rank-width and well-quasi-ordering of skew-symmetric or symmetric matrices

Sang-il Oum, KAIST, Daejeon, Korea

Suppose that a $V \times V$ matrix M has the following form:

$$M = \begin{array}{c} \\ Y \\ V \setminus Y \end{array} \begin{array}{c} \begin{array}{cc} Y & V \setminus Y \end{array} \\ \left(\begin{array}{cc} A & B \\ C & D \end{array} \right) \end{array}.$$

If $A = M[Y]$ is nonsingular, then we define the *Schur complement* (M/A) *of* A *in* M to be

$$(M/A) = D - CA^{-1}B.$$

We prove that every infinite sequence of skew-symmetric or symmetric matrices M_1, M_2, ... over a fixed finite field must have a pair M_i, M_j $(i < j)$ such that M_i is isomorphic to a principal submatrix of the Schur complement of a nonsingular principal submatrix in M_j, if those matrices have bounded rank-width. This generalizes three theorems on well-quasi-ordering of graphs or matroids admitting good tree-like decompositions; (1) Robertson and Seymour's theorem for graphs of bounded tree-width, (2) Geelen, Gerards, and Whittle's theorem for matroids representable over a fixed finite field having bounded branch-width, and (3) Oum's theorem for graphs of bounded rank-width with respect to pivot-minors.

Spanning closed walks in 3-connected plane graphs

Kenta Ozeki, National Institute of Informatics, Japan

A well-known Tutte theorem states that every 4-connected planar graph has a Hamilton cycle, and there exist infinitely many 3-connected non-Hamiltonian plane graphs. Therefore, several researchers have tried to find some "good" structures which are close to Hamilton cycles in 3-connected plane graphs. For example, every 3-connected plane graph has a spanning closed walk which passes each vertex at most twice (by Gao and Richter [2]). In this paper, we concentrate on a spanning closed walk of few length. In fact, a hamilton cycle of a graph G of order n is a spanning closed walk of length n, so the shorter a spanning closed walk is, the closed to a hamilton cycle. Asano, Nishizeki and Watanabe showed that every triangulation G of the plane has a spanning closed walk of length at most $\max\left\{\frac{3}{2}(|G|-3), 0\right\}$. Recently we improved it as follows:

Theorem 1 *Let G be a 3-connected plane graph of order n. Then G has a spanning closed walk of length at most $\frac{4}{3}(n-1)$.*

This is a joint work with K. Kawarabayashi (National Institute of Informatics).

Similarly to the result by Gao and Richter [2], consider a spanning closed walk of short length which passes each vertex at most twice. This seems a good problem, but no one succeeded to give an answer to it. Actually, we posed the following conjecture.

Conjecture 1 *Let G be a 3-connected plane graph of order n. Then G has a spanning closed walk of length at most $\frac{4}{3}(n-1)$ which passes each vertex at most twice.*

References

[1] T. Asano, T. Nishizeki and T. Watanabe, An upper bound on the length of a Hamiltonian walk of a maximal planar graph, J. Graph Theory **4** (1980) 315–336.

[2] Z. Gao and R.B. Richter, 2-walks in circuit graphs, J. Combin Theory Ser. B **62** (1994) 259–267.

Weighted disjoint paths

Gyula Pap, Egervary Research Group, Budapest, Hungary

Motivated by famous results of Mader and Karzanov, and recent progress made by Hirai, I started working on weighted disjoint paths problems. One can propose several nice problems related to disjoint paths, most of which being NP-hard. However, there are some surprising positive results, like Mader's disjoint paths problem, min cost flow, and the existence of a half-integral, or otherwise boundedly fractional optima for certain multiflow problems.

One question is to try to generalize Mader's disjoint S-paths theorem to a weighted setting, which is especially challenging in the node-disjoint case. In this general problem, we are given a graph, a subset S of nodes, and for every pair of nodes $s, t \in S$ we are given weight $w(s, t)$. The weight of an s, t-path P is determined by the weight of the pair of its endpoints, i.e. $w(P) = w(s, t)$. The goal then is to find a set of node-disjoint S-paths $P_1, P_2, \cdots$ that maximizes the sum of path weights $\sum_i w(P_i)$. Clearly, this general setting includes the disjoint S-paths problem that is solved by Mader. My main result here (yet to be published) is that for certain weight functions – so-called truncated tree metrics – this weighted disjoint paths problem is tractable, while for every weight function that is not a truncated tree metric, the problem becomes NP-hard. Thus we obtain a complete classification of weight functions for which the weighted disjoint path problem is tractable. The hardness result is proved by characterizing truncated tree metrics by a four-point condition, and then establishing a reduction to a integer 2-commodity flow problem. The tractability in case of truncated tree metric weights is shown by a combinatorial min-max formula, which is proved by a primal-dual algorithm using Mader's original theorem to determine a dual change. As a by-product, one can prove that some extended LP formulation has an integer optimum, given that the weights are equal to a truncated tree metric. It would be nice to prove the integrality of this mysterious LP directly.

With Hiroshi Hirai, we have also been working on the edge-disjoint case, where we proved a similar characterization of weight functions, with a quite different polyhedral approach.

Approximation Algorithms for Correlated Knaspacks and Non-Martingale Bandits

R. Ravi, RIMS and Tepper School of Business, CMU

In the stochastic knapsack problem, we are given a knapsack of size B, and a set of jobs whose sizes and rewards are drawn from a known probability distribution. However, the only way to know the actual size and reward is to schedule the job—when it completes, we get to know these values. How should we schedule jobs to maximize the expected total reward? We know constant-factor approximations for this problem when we assume that rewards and sizes are independent random variables, and that we cannot prematurely cancel jobs after we schedule them. What can we say when either or both of these assumptions are dropped?

Not only is the stochastic knapsack problem of interest in its own right, but techniques developed for it are applicable to other stochastic packing problems. Indeed, ideas for this problem have been useful for budgeted learning problems, where one is given several arms which evolve in a specified stochastic fashion with each pull, and the goal is to pull the arms a total of B times to maximize the reward obtained. Much recent work on this problem focus on the case when the evolution of the arms follows a martingale, i.e., when the expected reward from the future is the same as the reward at the current state. However, what can we say when the rewards do not form a martingale?

We give constant-factor approximation algorithms for the stochastic knapsack problem with correlations and cancellations, and also for some budgeted learning problems where the martingale condition is not satisfied, using similar ideas. Indeed, we can show that previously proposed linear programming relaxations for these problems have large integrality gaps. We propose new time-indexed LP relaxations; using a decomposition and "shifting" approach, we convert these fractional solutions to distributions over strategies, and then use the LP values and the time ordering information from these strategies to devise a randomized scheduling algorithm. We hope our LP formulation and decomposition methods may provide a new way to address other correlated bandit problems with more general contexts. We are currently working on extensions of these methods to stochastic Orienteering problems.

This is joint work with Anupam Gupta, Ravishankar Krishnaswamy and Marco Molinaro, all from CMU.

Collapse of the VPN Pyramid

András Sebő, CNRS, Grenoble

The celebrated VPN Tree Conjecture raised by Fingerhut, Suri and Turner and then again by A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yenerhas has been proved by N. Goyal, N. Olver, and B. Shepherd using an intermediate station, the equivalence of the so called "Pyramidal Conjecture" by

F. Grandoni, V. Kaibel, G. Oriolo, and M. Skutella (later a shortcut has been observed using a result of Padberg and Rao).

The problem consists in designing paths between a given set of terminals so as to minimize the cost of capacities to be bought when routing a demand function – satisfying certain linear constraints – between terminals through the designed paths. Until now the results have been built as a pyramid using bricks from previous work as black boxes (or black bricks). In this note we redigest the subject with a geometric insight that leads to a simpler proof and sharpening the results. The black boxes are opened, the proof pyramid collapses to determining the extreme points of a polytope. (Incidentally he "pyramidal conjecture" turns out to lose its pyramidal character.)

It turns out that the problem has a natural formulation as optimizing on a convex set whose vertices correspond to Steiner-trees of the network. Besides the simple proof and the connexion to polyhedral combinatorics, this new look on the problem yields - besides the known polynomial algorithms - to simple proofs and sharper results. Joint work with Nicola Apollonio, Fabrizio Grandoni, and Gianpaolo Oriolo.

My main interest :

I am interested in packing, covering, both combinatorial and polyhedral aspects, related integer progamming problems or their generalizations formulated as geometry of numbers problems; various algorithmic or structural questions related to combinatorial optimization problems. Let me mention three concrete problems I am alernately working on these days :

VPN : This is a famous problem of computer science that has been studied in the past years. With my colleagues (see the abstract below) we initiated a polyhedral study that led to simpler proofs abd a deeper understanding, and sharpening the results. See more details in the abstract of a suggested talk below.

ADDITIVE GAPS : The difference between optimal primal and dual integer solutions of linear programs, or the difference between the optimal integer solution and the linear programming optimum.

The integer rounding property is a classical example when this gap is small. Recently some new examples and operations (from different authors including me) and problerms occurred. Even when the integer rounding property does not hold, for certain general classes of interesting problems this gap is surprisingly small : Vizing's theorem for and the Goldberg-Seymour conjecture for the chromatic index are classical examples. (They can be viewed as problems on the gap of linear programs concerning matchings.)

I am interested - with Gennady Shmonin - in the bin packing problem, where the gap of an appropriate formulation might be 1. With András Gyárfás and Nicolas Trotignon we were exploring the chromatic gap, the difference between the chromatic number and the clique number which involves both matching theory and Ramsey-theory.We have submitted a paper on our results. We are studying now how the difference (or other relation) between the chromatic number and the clique-number can be bounded if we have bounds for this difference in the complement of a graph. Longstanding conjectures of Gyárfás summarize the main challenges in this area. (I spoke a lot about these subjects last year, I am fed up and some of the participants even more ;=))

PACKING AND COVERING : This is an ongoing work with Matěj Stehlík.

Starting with an easily stated hypergraph reformulation of matching problems in graphs we get a key for a simple rewriting of some results on packing and covering, and turning towards new ones. With our new look on this problem, a difficult theorem of Gallai on color-critical graphs is equivalent to his own celebrated lemma on "matching-stable" graphs being factor-critical. Using this, we are revisiting other results and problems from the subject of extremal combinatorics and blocking, antiblocking. (It would be too early to speak about this.)

Polynomial-Time Approximation Scheme for Maximizing Gross Substitutes Utility under Budget Constraints

Akiyoshi Shioura, Tohoku University

We consider the problem of maximizing a gross substitutes utility function under a constant number of budget (knapsack) constraints. This problem often appears in mathematical economics and (algorithmic) game theory. We show that there exists a polynomial-time approximation scheme (PTAS) for this problem.

Mathematically, a gross substitutes utility function is defined as a set function $f : 2^N \to \mathbb{R}$, where N is a finite set, satisfying the following condition:

$\forall p \in \mathbb{R}^N$, $\forall X \in \arg\max_{U \subseteq N}\{f(U) - p(U)\}$, $\forall \alpha \in \mathbb{R}_+$, $\forall i \in N$,
$\exists Y \in \arg\max_{U \subseteq N}\{f(U) - (p + \alpha\chi_i)(U)\}$ such that $X \setminus \{i\} \subseteq Y$.

The concept of gross substitutes utility function is introduced in Kelso and Crawford (1982), where the existence of core and equilibrium is shown in a fairly general two-sided matching model. Since then, this condition has become a benchmark that is widely used in matching, auction, housing, and labor market models.

It is shown by Fujishige and Yang (2003) that a gross substitutes utility function is a subclass of M$^\natural$-concave functions, which is introduced by Murota and Shioura (1999) in the context of discrete convex analysis. A set function $f : \mathcal{F} \to \mathbb{R}$ defined on $\mathcal{F} \subseteq 2^N$ is said to be M$^\natural$-concave if it satisfies the following condition:

$\forall X, Y \in \mathcal{F}$, $\forall u \in X \setminus Y$, either (i) or (ii) (or both) holds:
(i) $X - u, Y + u \in \mathcal{F}$ and $f(X) + f(Y) \leq f(X - u) + f(Y + u)$,
(ii) $\exists v \in Y \setminus X$: $X - u + v$, $Y + u - v \in \mathcal{F}$ and $f(X) + f(Y) \leq f(X - u + v) + f(Y + u - v)$.

It is shown in Fujishige and Yang (2003) that a gross substitutes utility function is nothing but an M$^\natural$-concave function with $\mathcal{F} = 2^N$.

Based on this fact, we consider a more general problem of maximizing an M$^\natural$-concave function under budget constraints. This generalized problem includes, as a very special case, the problem of maximizing a linear function subject to a single matroid constraints and budget constraints discussed in Grandoni and Zenklusen (2010). We show that the generalized problem has a deterministic PTAS.

Our PTAS is obtained by extending the approach in Grandoni and Zenklusen (2010) for the budgeted linear function maximization. The extension is, however not straightforward since our problem deals with nonlinear objective functions. The approach consists of the following three major steps, combined with a partial enumeration:

(i) Construct a continuous relaxation of the original problem,
(ii) Compute an optimal fractional solution of the continuous relaxation,
(iii) Round the optimal fractional solution to a feasible solution of the original problem.

We construct a continuous relaxation of our problem by using the fact that an $M^\natural$-concave function can be extended to an ordinary concave function. Therefore, the continuous relaxation of our problem is a nonlinear programming problem in continuous variables. It is, however, not clear how to solve the continuous relaxation in polynomial time. The concave extension of $M^\natural$-concave function has no analytical formula, and is represented as the minimum of infinitely many linear functions. This makes it quite difficult to solve the continuous relaxation efficiently. We overcome this difficulty by making full use of the conjugacy results of $M^\natural$-concave functions in the theory of discrete convex analysis. We define a "vertex" optimal solution of the continuous relaxation, and show that such an optimal solution can be computed in polynomial time. Finally, we show that a vertex optimal solution can be rounded to a feasible solution by a simple rounding at the cost of decreasing the function value slightly.

A Randomized Rounding Approach for Symmetric TSP

Mohit Singh, McGill University

We show a $(3/2 - \epsilon)$-approximation algorithm for the graphical traveling salesman problem where the goal is to find a shortest tour in an unweighted graph G. This is a special case of the metric traveling salesman problem when the underlying metric is defined by shortest path distances in G. The result improves on the 3/2-approximation algorithm due to Christofides.

Similar to Christofides, our algorithm first finds a spanning tree whose cost is upper-bounded by the optimum and then it finds the minimum cost Eulerian augmentation of that tree. The main difference is in the selection of the spanning tree. Except in certain cases where the solution of LP is nearly integral, we select the spanning tree randomly by sampling from a maximum entropy distribution defined by the linear programming relaxation.

Despite the simplicity of the algorithm, the analysis builds on a variety of ideas such as properties of strong Rayleigh measures from probability theory, graph theoretical results on the structure of near minimum cuts, and the integrality of the T-join polytope from polyhedral theory. This is joint work with Shayan Oveis Gharan and Amin Saberi.

The properties of maximum entropy distribution defined by the LP solution had been a crucial ingredient in the recent $O(\frac{\log n}{\log \log n})$-approximation for asymmetric TSP given by Asadpour et al [1] and was also a motivation for our work. Two natural open questions thus arise.

1. Can the results be extended to all metrics to improve Christofides algorithm for metric TSP?

2. Are there other problems where maximum entropy distributions give improved approximation algorithms?

References

[1] Arash Asadpour, Michel X. Goemans, Aleksander Madry, Oveis Gharan, Shayan and Amin Saberi, *O(logn/loglogn) approximation to the asymmetric traveling salesman problem*, Proceeding of 21st ACM Symposium on Discrete Algorithms, 2010.

Exact and Approximate Shortest-Path Queries for Planar Graphs

Christian Sommer, MIT

Travel agencies or producers of navigation systems may wish to provide advice to clients, who want to know the shortest, fastest, or cheapest way from one point to another. Instead of searching a large part of a transportation map using a traditional algorithm (say Dijkstra's algorithm) at every client query, they could instead precompute certain information in order to better support subsequent queries. We design, analyze, and implement algorithms to precompute this information so as to efficiently obtain answers for point-to-point shortest path queries. The precomputed data structures are also called *distance oracles*.

We prove that it is possible to preprocess a planar graph on n nodes in time $O(n \log^2 n)$ into a data structure of size $O(n)$ such that $(1 + \epsilon)$–approximate distance queries (the resulting path is at most $(1 + \epsilon)$ times longer than the actual shortest path) can be answered in time $O(\epsilon^{-2} \log^2 n)$ for any $\epsilon > 0$ and exact distance queries can be answered in time $O(n^{0.501})$.

For exact queries, we can further improve the query time by increasing the space to $O(S)$ for a space parameter $S \in [n, n^2]$; then, the query time is at most $O(n/\sqrt{S})$ up to logarithmic factors. For example, the designer of a mobile navigation device, restricted by space constraints, can configure the trade off between space and query time such that the remaining space is well used to speed up user queries.

We also prove that we can preprocess a planar graph with tree-width w such that exact distance queries can be answered in time $O(w \log^2 n)$. As a consequence, we give a data structure using quasi-linear space that can answer exact shortest-path queries in time proportional to the number of edges on the shortest path (up to logarithmic factors). In recent years, some route-planning methods engineered by practitioners have been claimed to have this performance based on staggering experimental results. We can finally give an algorithm with guaranteed query time almost proportional to the path length.

This is joint work with Ken-ichi Kawarabayashi and Philip N. Klein (results on approximate distances), and Shay Mozes (results on exact distances).

An algorithmic decomposition of claw-free graphs leading to an $O(n^3)$-algorithm for the weighted stable set problem

Gautier Stauffer, INRIA - Bordeaux Institute of Mathematics

In a recent paper with Yuri Faenza and Gianpaolo Oriolo, we have proposed an algorithm for solving the maximum weighted stable set problem on claw-free graphs that runs in $O(n^3)$-time, drastically improving the previous best known complexity bound. This algorithm is based on a novel decomposition theorem for claw-free graphs, which is also introduced in the paper. We propose to present this new decomposition result by drawing a parallel with Krausz's (1943) characterization of line graphs. Despite being weaker than the recent tremendous structure result for claw-free graphs given by Chudnovsky and Seymour (2005–), our decomposition theorem is, on the other hand, algorithmic, i.e. it is coupled with an $O(n^3)$-time procedure that actually produces the decomposition. We believe that our algorithmic decomposition result is interesting on its own and might be also useful to solve other kind of problems on claw-free graphs.

An Excluded Minor Characterization of Seymour Graphs

Zoltán Szigeti, Laboratoire G-SCOP/Grenbole, France

A graph G is said to be a *Seymour graph* if for any edge set F there exist $|F|$ pairwise disjoint cuts each containing exactly one element of F, provided for every circuit C of G the necessary condition $|C \cap F| \leq |C \setminus F|$ is satisfied.

A first coNP characterization of Seymour graphs has been shown by Ageev, Kostochka and Szigeti, the recognition problem has been solved in a particular case by Gerards, and the related cut packing problem has been solved in the corresponding special cases. In this talk, we show new minor-producing operations that keep this property, and prove excluded minor characterizations of Seymour graphs: the operations are the contraction of full stars and that of odd circuits. This sharpens the previous results, providing at the same time a simpler and self-contained algorithmic proof of the existing characterizations as well, still using methods of matching theory and its generalizations.

Joint work with Alexander Ageev, Yohann Benchetrit, András Sebő

Computing the Maximum Degree of Minors in Mixed Polynomial Matrices via Combinatorial Relaxation

Mizuyo Takamatsu, Faculty of Science and Engineering, Chuo University

We present an algorithm for computing the maximum degree of minors in mixed polynomial matrices. Mixed polynomial matrices are polynomial matrices with two kinds of nonzero coefficients: fixed constants that account for conservation laws and independent parameters that represent physical characteristics. The computation of their maximum degrees of minors is known to be reduced to valuated independent assignment problems, which can be solved by polynomial numbers of additions, subtractions, and multiplications of rational functions. However, these arithmetic operations on rational functions are much more expensive than those on constants.

Our algorithm is based on the framework of combinatorial relaxation. In the algorithm, we find a combinatorial estimate of the solution by solving a weighted bipartite matching problem, and check if the estimate is equal to the solution by solving an independent matching problem. The algorithm mainly relies on fast combinatorial algorithms and performs numerical computation only when necessary. In addition, it requires no arithmetic operations on rational functions. As a byproduct, this method yields a new algorithm for solving a linear valuated independent assignment problem.

This is a joint work with Satoru Iwata and will appear at IPCO 2011.

Optimal matching forests and valuated delta-matroids

Kenjiro Takazawa, RIMS, Kyoto University

The matching forest problem in mixed graphs is a common generalization of the matching problem in undirected graphs and the branching problem in directed graphs. Giles (1982) presented an $\mathrm{O}(n^2m)$-time algorithm for finding a maximum-weight matching forest, where n is the number of vertices and m is that of edges, and a linear system describing the matching forest polytope. Later, Schrijver (2000) proved total dual integrality of the linear system.

In this talk, we reveal another nice property of matching forests: the degree sequences of the matching forests in any mixed graph form a delta-matroid and the weighted matching forests induce a valuated delta-matroid. We remark that the delta-matroid is not necessaliry even, and the valuated delta-matroid induced by weighted matching forests slightly generalizes the well-known notion of Dress and Wenzel's valuated delta-matroids. By focusing on the delta-matroid structure and reviewing Giles' algorithm, we design a simpler $\mathrm{O}(n^2m)$-time algorithm for the weighted matching forest problem. We also present a faster $\mathrm{O}(n^3)$-time algorithm by using Gabow's method (1973) for the weighted matching problem.

Generic Rigidity Matroids and the Dilworth Truncation

Shin-ichi Tanigawa, Kyoto University

One of the main topics in rigidity theory is to find a good characterization of generic rigidity of bar-joint frameworks. After Laman's result on 2-dimensional generic rigidity in 1970, it is still an important unsolved problem to find the 3-dimensional counterpart. One approach toward a combinatorial characterization of bar-joint frameworks is to investigate a special class of structural models, and several partial results are known for, e.g., the body-bar model (Tay 1984), the body-hinge model (Whiteley 1988), the body-bar-hinge model (Jackson and Jordán 2007), and the rod-bar model (Tay 1989,1991). Currently, every solvable case is characterized by so-called Maxwell's condition, where the rank function of the corresponding generic rigidity matroid is written by the Dilworth truncation of some submodular function.

Lovász (1977) or Mason (1977) gave a geometric interpretation of the Dilworth truncation for linear polymatroids; roughly, it can be considered as an operation of restricting the corresponding flats (i.e., linear subspaces) to a "generic" hyperplane in a projective space. This interpretation was then applied to show an alternative proof of Laman's theorem (Lovász and Yemini 1982). Specifically, they proved that the 2-dimensional generic rigidity matroid is obtained from the union of two graphic matroids by the Dilworth truncation operation.

Inspired by the alternative proof of Laman's theorem by Lovász and Yemini, we prove that the linear matroid that defines the generic rigidity of d-dimensional body-rod-bar frameworks can be obtained from the union of $\binom{d+1}{2}$ graphic matroids by applying variants of Dilworth truncation operations n_r times, where n_r denotes the number of rods. (Here, a body-rod-bar framework is a structure consisting of disjoint bodies and rods mutually linked by bars.) In these operations, each hyperplane is inserted so that it cuts only a specified subspace, and we are able to show that a naturally extended statement of the Dilworth truncation is true for the resulting set of flats even though a hyperplane is not generic. I am wondering when the restriction to such a "non-generic" hyperplane leads to the rank formula in the form of the Dilworth truncation.

Another interesting example can be found in the work by Servatius and Whiteley (1999) on the generic rigidity of length-and-direction mixed framework. This result implies that, for disjoint sets $\mathcal{F}_1$ and $\mathcal{F}_2$ of flats representing the union of two graphic matroids, the simultaneous restriction of $\mathcal{F}_1$ to H_1 and $\mathcal{F}_2$ to H_2 leads to the rank formula of $(\mathcal{F}_1 \cap H_1) \cup (\mathcal{F}_2 \cap H_2)$ in the form of the Dilworth truncation even if H_1 and H_2 are mutually related (although their proof is based on the different approach).

Efficient Enumeration

Takeaki Uno, National Institute of Informatics

My recent research interest is on enumeration algorithms. Compared to optimization, enumeration has not been studied actively. However, in some areas in informatics such as information retrieval, data mining and data engineering, enumerational approaches are often chosen, because of the uncertainties of objective functions, variety of usages, constructing noise-tolerant methods, and so on. Enumeration is often considered as a part of optimization, but the key techniques are often different; characterization of optimal solutions is key to optimization, but good neighbor relations for efficient search strategy is important for enumeration. One of my research goal is to clarify the difficulty of enumeration.

For recent large scale data, non-algorithmic improvements of enumeration algorithms are usually not enough to make the computation time practically short. Hypergraph dualization is a problem to find all minimal subsets of a set E intersecting any member of the given set family $\mathcal{F}$. This is a problem of enumerating minimal elements, and in the case that the sizes of minimal elements are very small (quite usual in real world data), minimal element enumeration is usually difficult compared to maximal element enumeration. We constructed a set system such that each minimal element of the original problem is a maximal element of the system (but not always), so that we can use existing maximal element enumeration techniques. We used simple properties to define the set system, thus we could speed up the iteration. By combining the recursive structure of the enumeration and the improvement, we could drastically reduce the computation time, so that our algorithm often terminates in few seconds while existing algorithms do not terminate in one hour. The enumeration of infrequent patterns, association rules, explanations, and anonymization are also minimal element enumeration. Characterization of minimal elements and developing efficient algorithms for these problems are interesting future works.

Connectivity augmentation and matching problems

László Végh, Eötvös University / Georgia Tech

In the node-connectivity augmentation problem, we want to add a minimum number of new edges to an undirected graph to make it k-node-connected. The complexity of this question is still open, although the analogous questions of both directed and undirected edge-connectivity and directed node-connectivity augmentation are known to be polynomially solvable.

We give a min-max formula and a polynomial time algorithm for the special case when the input graph is already $(k-1)$-connected, as conjectured by Frank and Jordan in 1994.

Towards the general problem, an important observation is that connectivity augmentation problems are equivalent to certain matching problems in the complement graph. Making a graph on n nodes $(n-2)$-connected is equivalent to finding a maximum matching in the complement. For $k = n-3$, we get the long-standing open question of finding a maximum square-free 2-matching.

Separators in Minor-free Graphs with Applications

Christian Wulff-Nilsen, School of Computer Science, Carleton University

Graph separators have proven to be a very useful tool in the development of efficient algorithms for several fundamental graph problems. For planar graphs, the separator theorems of Lipton and Tarjan and of Miller are at the heart of the currently fastest algorithms for e.g. shortest paths and min cuts. Lately, I have considered the problem of computing separators for a larger class of graphs, namely those that exclude a fixed minor. Alon, Seymour, and Thomas generalized the planar separator theorem to this class but their algorithm is slower than that of Lipton and Tarjan. A new separator algorithm was recently presented by Kawarabayashi and Reed. Its running time as a function of the size n of the input graph is better than that of Alon, Seymour, and Thomas but it comes at a cost: in addition to being very complicated, this new algorithm has a time dependency on the size h of the excluded minor which is huge, namely a tower function of h whose height is itself a function of h. My current research focuses on speeding up the algorithm of Alon, Seymour, and Thomas as a function of n while keeping a (small) polynomial dependency on h. Another result that I have been working on is a speed-up of Yuster's SSSP algorithm for minor-free graphs with negative edge weights. The speed-up comes from a new application of the separator theorem of Reed and Wood. A very recent area of interest to me is dynamic algorithms, in particular for the subgraph connectivity problem. I wish to explore this field and hopefully find improved algorithms for general graphs and/or for graphs with small separators.

Participants

Zdeněk Dvořák, rakdver@kam.mff.cuni.cz
Tamás Fleiner, fleiner@cs.elte.hu
Satoru Fujishige, fujishig@kurims.kyoto-u.ac.jp
Takuro Fukunaga, takuro@amp.i.kyoto-u.ac.jp
Mohammad Hajiaghayi, hajiagha@MIT.EDU
Hiroshi Hirai, hirai@kurims.kyoto-u.ac.jp
Petr Hliněný, hlineny@fi.muni.cz
Satoru Iwata, iwata@kurims.kyoto-u.ac.jp
Tibor Jordán, jordan@cs.elte.hu
Naonori Kakimura, kakimura@mist.i.u-tokyo.ac.jp
Naoyuki Kamiyama, kamiyama@ise.chuo-u.ac.jp
Ken-ichi Kawarabayashi, k_keniti@nii.ac.jp
Shuji Kijima, kijima@inf.kyushu-u.ac.jp
Zoltán Király, kiraly@cs.elte.hu
Philip Klein, klein@cs.brown.edu
Yusuke Kobayashi, kobayashi@mist.i.u-tokyo.ac.jp
Shungo Koichi, shungo@nanzan-u.ac.jp
Daniel Král', kral@kam.mff.cuni.cz
Lap Chi Lau, chi@cse.cuhk.edu.hk
Tom McCormick, Tom.McCormick@sauder.ubc.ca
Kazuo Murota, murota@mist.i.u-tokyo.ac.jp
Kiyohito Nagano, nagano@is.titech.ac.jp
Gianpaolo Oriolo, oriolo@disp.uniroma2.it
Sang-il Oum, sangil@kaist.edu
Kenta Ozeki, ozeki@nii.ac.jp
Gyula Pap, gyuszko2000@yahoo.com
R. Ravi, ravi@qatar.cmu.edu
András Sebő, andras.sebo@g-scop.inpg.fr
Akiyoshi Shioura, shioura@dais.is.tohoku.ac.jp
Mohit Singh, mohit@cs.mcgill.ca
Christian Sommer, sommer@is.s.u-tokyo.ac.jp
Gautier Stauffer, gautier.stauffer@math.u-bordeaux1.fr
Zoltán Szigeti, zoltan.szigeti@g-scop.inpg.fr
Mizuyo Takamatsu, takamatsu@ise.chuo-u.ac.jp
Kenjiro Takazawa, takazawa@kurims.kyoto-u.ac.jp
Shin-ichi Tanigawa, tanigawa@kurims.kyoto-u.ac.jp
Takeaki Uno, uno@nii.ac.jp
László Végh, veghal@cs.elte.hu
Christian Wulff-Nilsen, koolooz@diku.dk