

OPTIMA 90

Mathematical Optimization Society Newsletter

Philippe L. Toint

MOS Chair's Column

December 1, 2012. I am certain that all of you feel the same way as I do about the Berlin ISMP 2012: what a fantastic meeting! Among many highlights, let me (very subjectively) select a few.

The first was the opening ceremony. Not that I found the role I played great, but I had the clear feeling that the mathematical programming community was pleased (or even happy) to be gathered again, much like at a (very big) family event. The general mood was clearly positive and the enthusiasm was, at least from my point of view, nearly palpable among the prize festivities (see Katya's report on this in this issue) ...

The second thing I really liked was, obviously, the scientific content of the meeting. I found the choice of plenaries really good and of sufficiently broad interest to enjoy large audiences every time. The idea of the historical talks was also truly excellent and many colleagues mentioned to me how pleased they felt about it. I will not dwell on the details of the technical sessions I went to, because this is too specialized, but, like everyone else, I enjoyed learning what our colleagues (and often friends) were up to. Some great moments ...

The conference dinner met all my expectations. When Martin Skutella and I discussed the form of this event during the preparation months, we both agreed that it had to be an open affair, affordable for all and generally informal. This indeed was the case, but was also combined with an excellent setting (emphasized by the good weather, I have to say). And for those who took the boat trip on the Spree, this was really a great way to discover the architectural facets of history-rich Berlin. Again a memorable evening!

Talking of Berlin, there is no doubt that the city provided a most remarkable setting for our ISMP. I admit that my visit this time was by far the best as far as the impression of the city is concerned, probably helped, like the conference dinner, by the exceptional weather. Martin and his colleagues were able to order for us. When I left after the very nice farewell gathering, I promised myself I would go back quickly (I am returning this month).

Finally, the ISMP is most certainly the high point of a MOS chairman's term. As such, there is an inevitable amount of stress associated with it: will it be successful? Will the science and the organization live up to the usual high standards? In my case, I am

very much indebted to Martin and his great team for relieving me entirely of this worry. Not only the preparations gave no sign of grave difficulties, but the event itself was superbly managed, with efficiency and a large friendly smile. Many thanks again, dear Martin; I will definitely remember the Berlin ISMP days as amongst the best in my career.

Of course, I wish the same to the MOS chairman elect, Bill Cook, who will be taking over from me at the end of August 2013, and who will supervise the preparations of ISMP 2015 in Pittsburgh, with the help of François Margot. I know the preparations have started already, and I am surely looking forward to another great conference (and, even better, with a new chairman in charge)!

In the mean time, winter is arriving in Europe, days are getting short and windy, with the famous Belgian rain all too often present. A good time to start enjoying doing some mathematics in a corner by a wood fire ... I wish you all a happy and productive end of 2012.

Note from the Editors

In this issue you'll find an overview by Jens Vygen of the fascinating approximation results for the Traveling Salesman Problem that have recently been obtained by several groups after advances in the treatment of these questions had been stalling for many years. This survey article is accompanied by Michel Goeman's discussion column, in which he in particular reveals the problem he likes to see solved until his (hopefully far away) retirement.

Sam Burer,
Volker Kaibel
and Katya Scheinberg
Optima editors

Jens Vygen

New approximation algorithms for the TSP

The traveling salesman problem (TSP) is probably the best known combinatorial optimization problem. Although studied intensively for sixty years, the TSP continues to pose grand challenges. Cook's recent book (19) gives an excellent introduction.

Since the TSP is NP -hard ((46)), it is natural to ask for approximation algorithms. How good solutions can we guarantee to find in polynomial time? (18) devised a $\frac{3}{2}$ -approximation algorithm for the SYMMETRIC TSP: it always finds a solution that is at most 50% longer than optimum.

Can we do better? Can we do similarly well for the ASYMMETRIC TSP? These questions – still unsolved – belong to the most intriguing open problems in our field. Recently, there has been progress that makes us hope that we will learn more in the near future.

Contents of Issue 90 / December 2012

- 1 Philippe L. Toint, *MOS Chair's Column*
- 1 Note from the Editors
- 1 Jens Vygen, *New approximation algorithms for the TSP*
- 12 Michel X. Goemans, *Thinness Spurs Progress*
- 14 Prizes awarded at ISMP 2012
- 16 IPCO 2013
- 16 MIP 2013
- 16 Imprint

In this survey we try to describe the state of the art, in particular the recent progress, mostly from 2010–2012. We use standard notation and some basic terminology and well-known results from combinatorial optimization; see (49) or (64) if necessary.

1 Introduction

Let us first review different formulations of the problems and basic approximation algorithms.

1.1 ASYMMETRIC TSP and SYMMETRIC TSP

The ASYMMETRIC TSP can be defined as follows. Given a finite set V (of cities) and distances $c(v, w) \geq 0$ for all $v, w \in V$ (also called *length* or *cost*), find a closed walk of minimum total length visiting each city at least once. More precisely, we look for a sequence v_0, v_1, \dots, v_k with $v_k = v_0$ and $\{v_0, \dots, v_k\} = V$ (also called a *tour*) such that $\sum_{i=1}^k c(v_{i-1}, v_i)$ is minimum.

The SYMMETRIC TSP is the special case in which the distances are symmetric: $c(v, w) = c(w, v)$ for all $v, w \in V$.

Often these problems are formulated such that each city must be visited *exactly* once instead of at least once (except, of course, that we must end in the same city where we start). This is equivalent if the distances obey the *triangle inequality*

$$c(u, w) \leq c(u, v) + c(v, w) \quad \text{for all } u, v, w \in V \quad (1)$$

because then we can shortcut whenever we visit a city a second time.

On the other hand, given an instance of the ASYMMETRIC TSP (or SYMMETRIC TSP) as defined above, we can set

$$\bar{c}(v, w) := \min\{\sum_{e \in E(P)} c(e) : P \text{ path from } v \text{ to } w, V(P) \subseteq V\}$$

and consider (V, \bar{c}) instead. Note that \bar{c} obeys (1). The instance (V, \bar{c}) is equivalent to (V, c) because we can move from v to w at cost $\bar{c}(v, w)$ via a shortest v - w -path. The pair (V, \bar{c}) is called the *metric closure* of (V, c) .

1.2 Approximation algorithms

A ρ -*approximation algorithm* (for a minimization problem) is an algorithm that runs in polynomial time and always computes a solution (here: a tour) that costs at most ρ times the optimum. Here ρ can be a constant or a function of n ; here and henceforth $n = |V|$ denotes the number of cities.

If we do not require the triangle inequality but still want to visit every city exactly once, the problems look hopeless: any approximation algorithm would allow us to decide in polynomial time whether a given graph contains a Hamiltonian circuit, and thus imply $P = NP$ (this easy observation is due to (62)).

So we assume henceforth that we may visit cities more than once.

1.3 Euler's theorem

For the total length of a tour, all that matters is how many times we move from v to w for each ordered pair (v, w) , or how many times we move between v and w for each unordered pair $\{v, w\}$ in the undirected case. Hence we can represent a tour by a directed or undirected graph (possibly with parallel edges) with vertex set V . As observed by (28), this graph has two properties:

- (a) it is connected;
- and for every city:
- (b') the in-degree equals the out-degree in the directed case;
- (b'') the degree is even in the undirected case.

Digraphs with the properties (a) and (b') and undirected graphs with properties (a) and (b'') are called *Eulerian*. The conditions are equivalent to the existence of an *Eulerian walk*: a closed walk traversing

every edge exactly once and every vertex at least once. Given an Eulerian graph or digraph, one can find an Eulerian walk in linear time ((41)).

Therefore, one can reformulate the ASYMMETRIC TSP and the SYMMETRIC TSP by asking for a (multi)set F such that (V, F) is an Eulerian (directed or undirected, respectively) graph with minimum total length $c(F)$. We call such an F a *tour*, too. Here and in the following we abbreviate $c(F) := \sum_{e \in F} c(e)$ and $c(e) := c(v, w)$ for any edge e from v to w .

1.4 GRAPHIC TSP

A natural special case of the SYMMETRIC TSP arises when we are given a connected undirected graph G and let $V = V(G)$ and $c(v, w) = 1$ if $\{v, w\} \in E(G)$ and $c(v, w) = \infty$ otherwise. This problem is called the GRAPHIC TSP. The metric closure (V, \bar{c}) of (V, c) is also called the metric closure of G . Functions \bar{c} arising in this way are called *graphic metrics*.

By the observations in Section 1.3, the GRAPHIC TSP can be reformulated as follows. Given a connected graph G , find an Eulerian spanning multi-subgraph (V, F) with minimum $|F|$. Here a multi-subgraph arises from a subgraph by doubling a subset of its edges. Again, we call an Eulerian multi-subgraph of G simply a *tour* in G .

The GRAPHIC TSP has also been called GRAPH TSP by some authors. It is easy to see that, without loss of generality, we may assume that G is 2-vertex-connected (otherwise find a tour in each block separately).

1.5 Double tree and Christofides' algorithm

A 2-approximation algorithm for the SYMMETRIC TSP is easy: take a tree on vertex set V with minimum total edge length (it is well-known that such a *minimum spanning tree* can be found efficiently, e.g., by the greedy algorithm), and double all its edges. Since any tour is connected and thus contains a spanning tree, a minimum spanning tree cannot be longer than an optimum tour. Hence we have a 2-approximation algorithm.

(18) showed how to improve this. His algorithm also begins by computing a minimum spanning tree (V, S) . But then, to correct the parities, it adds a minimum weight T_S -join, where $T_S = \{v \in V : |\delta_S(v)| \text{ odd}\}$ is the set of odd-degree vertices of (V, S) . For a set $T \subseteq V$, a T -join is a subset F of edges such that $|\delta_F(v)|$ is odd for $v \in T$ and even for $v \in V \setminus T$. See Figure 1. Of course S itself is a T_S -join, but we can do better: since every tour contains two disjoint T_S -joins (color the edges of a tour red and blue, changing the color whenever visiting an element of T_S for the first time, and finally delete pairs of parallel edges of the same color), the minimum weight of a T_S -join is at most half the length of an optimum tour.

So we have a $\frac{3}{2}$ -approximation algorithm for the SYMMETRIC TSP. Its running time is $O(n^3)$, dominated by the subroutine to find a minimum weight T_S -join.

This bound on the approximation ratio of Christofides' algorithm is tight even for the GRAPHIC TSP: for a Hamiltonian graph (so a tour of length n exists) that contains a spanning tree all whose vertices have odd degree, if we take such a spanning tree, we end up with $\frac{3}{2}n - 1$ edges.




Figure 1. Christofides' algorithm. From left to right: an instance of the GRAPHIC TSP, a spanning tree (V, S) whose odd-degree vertices (elements of T_S) are shown as squares, a minimum T_S -join, and the resulting tour.

2 Relaxations

For NP-hard problems it is often useful to study relaxations that are easier to solve. For the TSP, there are several interesting relaxations.

As explained in Section 1.3, it is often useful to view a tour as a (multi)set F of edges. Now we associate a vector $x \in \mathbb{Z}_{\geq 0}^E$ with each tour, where E is the set of ordered or (in the symmetric case) unordered pairs of elements of V and x_e is the number of copies of e in F . Then the tour has length $c(F) = c(x) := \sum_{e \in E} c(e)x_e$. Given a vector $x \in \mathbb{R}_{\geq 0}^E$, the graph $(V, \{e \in E : x_e > 0\})$ is called the *support graph* of x . For any subset $E' \subseteq E$ we will write $x(E') := \sum_{e \in E'} x_e$.

2.1 Subtour LP

Let (V, c) be an instance of the SYMMETRIC TSP, $n = |V| \geq 3$, and $E = \binom{V}{2}$. The following LP, first formulated by (22), has often been called *subtour elimination LP* or simply *subtour LP* or *Held-Karp relaxation*:

$$\begin{aligned} \min \quad & c(x) \\ \text{subject to} \quad & x(\delta(U)) \geq 2 \quad (\emptyset \neq U \subset V) \\ & x(\delta(v)) = 2 \quad (v \in V) \\ & x_e \leq 1 \quad (e \in E) \\ & x_e \geq 0 \quad (e \in E) \end{aligned} \tag{2}$$

Note that the constraints $x_e \leq 1$ ($e \in E$) could be omitted as they are implied by the other constraints: for $e = \{u, v\}$ we have $2x_e = x(\delta(u)) + x(\delta(v)) - x(\delta(\{u, v\})) \leq 2 + 2 - 2$.

The set of feasible solutions of the subtour LP (2) is called the *subtour polytope*. The integral feasible solutions of (2) are exactly the incidence vectors of Hamiltonian circuits. Their convex hull is called the *TSP polytope*.

So (2) is a relaxation of the SYMMETRIC TSP if the triangle inequality holds. For a general instance, we can consider (2) for its metric closure.

This relaxation has been tightened by many classes of additional valid inequalities. It is also the basis of branch-and-cut algorithms (with exponential worst-case running time) that made impressive progress over the last four decades and have found optimum solutions to TSP instances with up to 85 900 cities; see (3).

2.2 Integrality ratio

Obviously, the integer solutions of (2) are exactly the incidence vectors of Hamiltonian circuits. However, this LP has no integral optimum solution in general. Figure 2 shows a well-known example. We have an infinite family of graphs G ; each is an instance of the GRAPHIC TSP. The subtour LP of the metric closure (V, c) has a unique optimum solution: $x_e = 1$ on the horizontal edges and $x_e = \frac{1}{2}$ on the six other edges. Its LP value is n . However, an optimum tour has length $\frac{4}{3}n - 2$.

The *integrality ratio* of a family of polytopes $(P \subseteq \mathbb{R}^{E_P})_{P \in \mathcal{P}}$ is the supremum of $\min\{c(x) : x \in P \cap \mathbb{Z}^{E_P}\} / \min\{c(x) : x \in P\}$ over all $P \in \mathcal{P}$ and all $c : E_P \rightarrow \mathbb{R}_{>0}$. Often the weight functions are restricted in the supremum, e.g. to metrics or to graphic metrics. The above family of examples shows that the integrality ratio of the family of subtour polytopes for graphic metrics (and hence for general metrics) is at least $\frac{4}{3}$. Worse examples are not known. The fact that the worst known examples are instances of the GRAPHIC TSP raised interest in this special case. See also Sections 2.5 and 7.2.




Figure 2. Examples showing a lower bound of $\frac{4}{3}$ on the integrality ratio of the subtour polytope

2.3 Spanning trees

The difficulty of the SYMMETRIC TSP lies in the combination of connectivity and parity requirements. If we require only connectivity, a minimum spanning tree does the job. (25) gave the following polyhedral description:

Proposition 1. *The convex hull of incidence vectors of trees with vertex set V and edges in E is the set of vectors $x \in \mathbb{R}^E$ with*

$$\begin{aligned} x(E) &= n - 1 \\ \sum_{e=\{v,w\} \in E: v,w \in U} x_e &\leq |U| - 1 \quad (\emptyset \neq U \subset V) \\ x_e &\geq 0 \quad (e \in E) \end{aligned} \tag{3}$$

This set is called the *spanning tree polytope* of the graph (V, E) . The following easy observation was made by (5), strengthening a result of (40):

Proposition 2. *If x is a feasible solution of (2), then $\frac{n-1}{n}x$ is in the relative interior of the spanning tree polytope of the support graph.*

Proof. We have $\frac{n-1}{n}x(E) = \frac{n-1}{2n} \sum_{v \in V} x(\delta(v)) = n - 1$ as well as $\sum_{e=\{v,w\} \in E: v,w \in U} \frac{n-1}{n}x_e = \frac{n-1}{2n} (\sum_{v \in U} x(\delta(v)) - x(\delta(U))) = \frac{n-1}{2n} (2|U| - x(\delta(U))) \leq \frac{n-1}{n} (|U| - 1)$ for any $\emptyset \neq U \subset V$. \square

2.4 T-joins

Now consider the parity aspect. (26) proved:

Proposition 3. *The minimum weight of a T-join in a graph (V, E) with weights $c \in \mathbb{R}_{\geq 0}^E$ and $T \subseteq V$ equals the optimum value of the LP:*

$$\begin{aligned} \min \quad & c(x) \\ \text{subject to} \quad & x(\delta(U)) \geq 1 \quad (U \subseteq V, |U \cap T| \text{ odd}) \\ & x_e \geq 0 \quad (e \in E) \end{aligned} \tag{4}$$

The cuts $\delta(U)$ with $|U \cap T|$ odd are called *T-cuts*.

For negative weights the LP (4) cannot be used directly. We will also need:

Proposition 4. *The convex hull of incidence vectors of T-joins in (V, E) is the set of vectors $x \in [0, 1]^E$ with*

$$|F| - x(F) + x(\delta(U) \setminus F) \geq 1 \quad (U \subseteq V, F \subseteq \delta(U), |U \cap T| + |F| \text{ odd}) \tag{5}$$

This is called the *T-join polytope* of (V, E) . A minimum weight T-join can be found in $O(n^3)$ time via weighted matching. See (64) or (49) for details and proofs of Propositions 1, 3, and 4.

2.5 Wolsey's analysis

(69) proved that Christofides' algorithm computes a tour of length at most $\frac{3}{2}L$, where L is the LP value of (2). In fact, this is easy to see from the above: By Proposition 2, the minimum weight of a spanning tree is at most L . By Proposition 3, the minimum weight of a T-join is at most $\frac{L}{2}$ for any $T \subseteq V$ with $|V|$ even.

This shows that the integrality ratio of (2) is at most $\frac{3}{2}$. No better upper bound is known in general.

2.6 Two-edge-connected spanning subgraphs

Every tour is 2-edge-connected, so a relaxation of the SYMMETRIC TSP is to find a minimum weight 2-edge-connected spanning subgraph (if the triangle inequality holds) or multi-subgraph. Unfortunately, these problems are also NP-hard (cf. Section 7.6).

Let $G = (V, E)$ be a 2-edge-connected undirected graph. Then the incidence vectors of the 2-edge-connected spanning subgraphs (2ECSS) of G are the integral feasible solutions of the following LP:

$$\begin{aligned} \min \quad & c(x) \\ \text{subject to} \quad & x(\delta(U)) \geq 2 \quad (\emptyset \neq U \subset V) \\ & x_e \leq 1 \quad (e \in E) \\ & x_e \geq 0 \quad (e \in E) \end{aligned} \tag{6}$$

This LP arises from (2) by omitting the equality constraints. If we allow using edges twice, the LP (6) can be simplified further by omitting the upper bounds:

$$\begin{aligned} \min \quad & c(x) \\ \text{subject to} \quad & x(\delta(U)) \geq 2 \quad (\emptyset \neq U \subset V) \\ & x_e \geq 0 \quad (e \in E) \end{aligned} \quad (7)$$

Cunningham (see (54)) and (39) observed:

Proposition 5. *If (V, E) is a complete graph, $|V| \geq 3$, and c obeys the triangle inequality, then the optimum values of (2), (6), and (7) are the same.*

Proof. Let x be a rational feasible solution of (7). Choose $k \in \mathbb{N}$ such that kx_e is an even integer for each $e \in E$. If there is a vertex $v \in V$ with $x(\delta(v)) > 2$, choose incident edges $e = \{v, w\}$ and $e' = \{v, w'\}$ with $x_e > 0$ and $x_{e'} > 0$, reduce x_e and $x_{e'}$ each by $\frac{1}{k}$ and increase $x_{\{w, w'\}}$ by $\frac{1}{k}$ while maintaining feasibility (the existence of two such edges e, e' follows from applying Lovász' [1976] splitting theorem to the Eulerian graph with kx_e copies of each edge e). Note that we maintain the property that $kx(\delta(v))$ is an even integer for all $v \in V$, so we end up with a feasible solution of (7) also satisfying $x(\delta(v)) = 2$ for all $v \in V$. Then also $x_e = \frac{1}{2}(x(\delta(v)) + x(\delta(w)) - x(\delta(\{v, w\}))) = \frac{1}{2}(2 + 2 - x(\delta(\{v, w\}))) \leq 1$ for all $e = \{v, w\} \in E$, so we have a feasible solution of (2). Due to the triangle inequality we never increased $c(x)$. \square

We also note:

Proposition 6. *If (V, E) is a 2-edge-connected graph and $c(e) = 1$ for all $e \in E$, then the optimum values of (6) and (7) are the same.*

Proof. Let x be an optimum solution of (7). Let $f = \{v, w\} \in E$ with $x_f > 1$. Call two vertices a and b close if $x(\delta(\{a\} \cup S) \setminus \{f\}) \geq 1$ for all $S \subseteq V \setminus \{a, b\}$. This is a transitive relation. If v and w are close, then we can reduce x_f to 1 and maintain feasibility. Otherwise each vertex is either close to v or close to w ; so there is an edge $f' = \{v', w'\}$ such that v and v' are close, w and w' are close, and $x_{f'} < 1$. Then increasing $x_{f'}$ to $\min\{1, x_{f'} + x_f - 1\}$ and decreasing x_f to 1 maintains feasibility. \square

The same holds for integral solutions: we never need to take two copies of any edge, except of course for bridges of G .

The constraints of (7) define facets of the *graphical traveling salesman polyhedron*: the convex hull of vectors $x \in \mathbb{Z}_{\geq 0}^E$ for which $x(\delta(U)) \in \{2, 4, 6, \dots\}$ for all $\emptyset \neq U \subset V$. This was studied by (20).

2.7 Asymmetric subtour LP

Let (V, c) be an instance of the ASYMMETRIC TSP with (1) and $E = \{(v, w) : v, w \in V, v \neq w\}$. The following is the natural analogon of the subtour LP in this case:

$$\begin{aligned} \min \quad & c(y) \\ \text{subject to} \quad & y(\delta^+(U)) \geq 1 \quad (\emptyset \neq U \subset V) \\ & y(\delta^+(v)) = y(\delta^-(v)) = 1 \quad (v \in V) \\ & y_e \geq 0 \quad (e \in E) \end{aligned} \quad (8)$$

Again, the integral feasible solutions to this LP are exactly the Hamiltonian circuits. Vectors $y \in \mathbb{R}_{\geq 0}^E$ with $y(\delta^+(v)) = y(\delta^-(v))$ for all $v \in V$ are called *circulations* in (V, E) .

From a feasible solution y to (8) one can obtain a feasible solution to (2) by setting $x_{\{v, w\}} := y_{(v, w)} + y_{(w, v)}$ for all $\{v, w\} \in \binom{V}{2}$.

2.8 Solving the linear programs

All linear programs above have exponentially many constraints, but they can all be solved in polynomial time; in fact an optimum basic solution can be found in polynomial time. One way to show this is via the equivalence of optimization and separation. The LPs for spanning trees and T -joins can be solved by combinatorial algorithms. For the LPs (2), (6), (7), and (8), there are straightforward polynomial-size extended formulations (by introducing flow variables and using the max-flow min-cut theorem), but combinatorial algorithms to solve these LPs are not known. (40), however, showed how to solve (2) fast approximately.

2.9 Optimum basic solutions

Any optimum basic solution x^* of any of the LPs (2), (6), and (7) has at most $2n - 3$ nonzero variables; in fact the subgraph of the support graph induced by U has at most $2|U| - 3$ edges for any $U \subseteq V$ with $|U| \geq 2$ (cf. (20) and (38)).

Any optimum basic solution x^* of the LP (8) has at most $3n - 4$ nonzero variables (and the subgraph of the support graph induced by U has at most $3|U| - 4$ edges for any $U \subseteq V$ with $|U| \geq 2$); this was also shown by (38).

The basic feasible solutions of (8) arise as the unique solutions of a linear equation system with all coefficients 0 or 1, so by Cramer's rule each of their components can be written as $\frac{a}{b}$ for integers a and $b \leq (3n - 4)!/2$. The same holds for (2), (6), and (7), even with $b \leq (2n - 3)!/2$.

2.10 Covering the cities by disjoint circuits

Another relaxation works both in the directed and undirected case. We ignore connectivity and look for a graph in which each city belongs to a circuit and the circuits are pairwise vertex-disjoint.

In the undirected case, such an edge set is called a *perfect 2-matching* because every vertex must have degree 2. A minimum weight perfect 2-matching can be found in polynomial time (this is essentially equivalent to nonbipartite weighted matching), but it did not prove useful in the design of approximation algorithms so far.

In the directed case, the analogous relaxation is even easier to solve. Here every vertex must have in-degree and out-degree 1. A minimum weight spanning subgraph with this property can easily be found by solving a bipartite weighted matching problem. This was used for the first nontrivial approximation algorithm for the ASYMMETRIC TSP, to be described next.

2.11 $O(\log n)$ -approximation for ASYMMETRIC TSP

For the ASYMMETRIC TSP no constant-factor approximation algorithm is known, so we will consider $f(n)$ -approximation algorithms, where $f(n)$ is a function of the number n of cities. It is trivial to give an n -approximation algorithm: order the cities arbitrarily, say $V = \{v_1, \dots, v_n\}$, and take a shortest v_n - v_1 -path and a shortest v_{i-1} - v_i -path for $i = 2, \dots, n$.

The first nontrivial approximation algorithm was found by (33). It assumes that the triangle inequality holds and works as follows.

Begin with $W := V$. Find a minimum weight subset F of edges with $|\delta_F^+(v)| = |\delta_F^-(v)| = 1$ for all $v \in W$ and $|\delta_F^+(v)| = |\delta_F^-(v)| = 0$ for all $v \in V \setminus W$ (cf. Section 2.10). Then pick one vertex from each connected component of (W, F) , and replace W by the set of these vertices. Iterate until (W, F) is connected.

The set of all edges chosen in this algorithm forms an Eulerian graph. Due to the triangle inequality, the total cost of edges picked in each iteration is at most the length of an optimum tour. Since $|W|$ decreases by a factor of two in each iteration, we are done after $\lceil \log_2 n \rceil$ iterations. Hence we have a $(\log_2 n)$ -approximation algorithm.

(8), (44), and (29) improved this by a constant factor.

3 Random sampling

It is quite natural to first take a spanning tree to guarantee connectivity, and then add a minimum cost set of edges in order to make the graph Eulerian. This idea (underlying Christofides' algorithm) works also in the directed case as we shall see soon (in the proof of Theorem 8).

A minimum spanning tree does often not give the best overall result. A certain kind of random sampling led to better approximation algorithms for the ASYMMETRIC TSP and the GRAPHIC TSP.

3.1 Thin trees

Asadpour, Goemans, Mađry, Oveis Gharan and Saberi [2010] were the first to obtain an $o(\log n)$ -approximation algorithm for the ASYMMETRIC TSP. Their algorithm is randomized. The main ingredient is the following result, which, interestingly, applies to an undirected instance:

Theorem 7. *There is a randomized polynomial-time algorithm which, given a feasible solution x of the LP (2), computes a spanning tree (V, S) such that with probability at least $\frac{1}{2}$ we have $c(S) \leq 2c(x)$ and $|\delta_S(U)| \leq \alpha x(\delta(U))$ for all $U \subseteq V$, where $\alpha = 4 \log n / \log \log n$.*

The last property is called α -thinness. By Proposition 2, $\frac{n-1}{n}x$ is a convex combination of spanning trees, i.e. $\frac{n-1}{n}x_e = \sum_{S \in \mathcal{S}: e \in S} p_S$ for all $e \in E$, where \mathcal{S} is the set of (edge sets of) spanning trees, $p_S \geq 0$ for all $S \in \mathcal{S}$ and $\sum_{S \in \mathcal{S}} p_S = 1$. Such an explicit convex combination can be obtained in polynomial time. If we pick each tree $S \in \mathcal{S}$ with probability p_S , the expected cost is $\sum_{e \in E} c(e)x_e$, and hence the cost at most is twice as much with probability at least $\frac{1}{2}$.

The difficulty is that such a random spanning tree will in general not be thin enough. Therefore (5) choose the probability distribution carefully, namely such that it maximizes the entropy $\sum_{S \in \mathcal{S}} p_S \log \frac{1}{p_S}$. Equivalently, $p_S = \gamma \prod_{e \in S} \lambda_e$ for all $S \in \mathcal{S}$, for suitable positive numbers γ and λ_e ($e \in E$). Such a distribution is also called λ -uniform.

(5) show how to sample trees efficiently from approximately this distribution. Moreover, they prove that the random variables indicating for each edge whether it is part of the selected tree are negatively correlated; then thinness is implied by a Chernoff bound together with the fact that in any graph there are less than $2n^{2\gamma}$ many γ -approximate minimum cuts, for any $\gamma \geq 1$ ((45)). See (5) for the details. (Alternatively, a thin tree can be obtained by the dependent randomized rounding approach of (15).)

3.2 The $O(\log n / \log \log n)$ -approximation algorithm

Using Theorem 7, the randomized $O(\log n / \log \log n)$ -approximation algorithm of (5) and its analysis can be described easily. We work in the metric closure, so c satisfies the triangle inequality.

First solve the LP relaxation (8) to obtain a vector y . Then get a solution x of (2) by setting $x_{\{v,w\}} := y_{(v,w)} + y_{(w,v)}$ for all $\{v,w\} \in \binom{V}{2}$. Note that $x_e = 0$ or $x_e \geq \frac{1}{(3n-4)!}$ for all $e \in \binom{V}{2}$ (cf. Section 2.9); so x can be stored with $O(n^2 \log n)$ bits.

Next apply Theorem 7 to obtain a spanning tree (V, S) . Orient the edges of this tree by replacing each $\{v,w\} \in S$ by the cheaper one of (v,w) and (w,v) . Setting $c'(\{v,w\}) := \min\{c(v,w), c(w,v)\}$, we get with probability at least $\frac{1}{2}$ that the resulting arc set R satisfies $c(R) = c'(S) \leq 2c'(x) \leq 2c(y)$ as well as $|\delta_R^-(U)| \leq |\delta_S(U)| \leq \alpha x(\delta(U)) = \alpha(y(\delta^-(U)) + y(\delta^+(U))) = 2\alpha y(\delta^+(U))$ for all $U \subseteq V$.

Finally apply the following Theorem 8 to R and y . With probability at least $\frac{1}{2}$ we obtain a tour of length at most $(2\alpha + 2)c(y)$.

Theorem 8. *Let (V, R) be a connected spanning subgraph of the complete digraph (V, E) , $y \in \mathbb{R}_{\geq 0}^E$, and $\beta > 0$ such that $|\delta_R^-(U)| \leq \beta y(\delta^+(U))$ for all $U \subseteq V(G)$. Then we can find a tour F with $c(F) \leq c(R) + \beta c(y)$ in polynomial time.*

Proof. Let $l(e) := 1$ for $e \in R$ and $l(e) := 0$ for $e \in E \setminus R$. Any integral circulation f in (V, E) with $f \geq l$ corresponds to a tour. We compute an integral minimum cost circulation $f^* \geq l$ and note that the resulting tour has cost $c(f^*)$.

To prove that such a circulation (and hence an integral circulation) of cost at most $c(R) + \beta c(y)$ exists, we let $u(e) := \max\{l(e), \beta y_e\}$ for all $e \in E$ and observe that a circulation g with $l \leq g \leq u$ exists; then $c(f^*) \leq c(g) \leq \sum_{e \in E} c(e)u(e) \leq c(R) + \beta c(y)$.

The existence of g follows from Hoffman's [1960] circulation theorem: we have $l \leq u$ and $l(\delta^-(U)) = |\delta_R^-(U)| \leq \beta y(\delta^+(U)) \leq u(\delta^+(U))$ for all $U \subseteq V$. \square

3.3 Random sampling for the SYMMETRIC TSP

The random sampling of (5) was also used by (56) for the first improvement over Christofides' algorithm for the GRAPHIC TSP.

They proposed the following algorithm. Take the metric closure and solve the subtour LP (2) to obtain an optimum basic solution x . Again, $\frac{n-1}{n}x$ is a convex combination of spanning trees, and we pick one at random according to a maximum entropy distribution; call it (V, S) . Let T_S be again the set of odd degree vertices of (V, S) . Finally add a minimum-weight T_S -join to S as in Christofides' algorithm.

(56) conjectured that this algorithm has better approximation ratio than $\frac{3}{2}$, but they could prove this only for the GRAPHIC TSP, and only for a slight variant of this algorithm. Their main structure theorem is the following. (The constants below are not best possible, but the improvement is tiny anyway.)

Theorem 9. *Let (V, c) be an instance of the SYMMETRIC TSP satisfying the triangle inequality. Let x be an optimum solution of (2), and let (V, S) be a spanning tree picked at random according to the maximum entropy distribution $(p_S)_{S \in \mathcal{S}}$ with $\sum_{S \in \mathcal{S}: e \in S} p_S = \frac{n-1}{n}x_e$ for all $e \in E$. Let T_S be the set of odd degree vertices of (V, S) . Call an edge e good if e does not belong to any T_S -cut $\delta(U)$ with $x(\delta(U)) \leq 2 + 10^{-15}$. Then at least one of the following holds:*

- (a) *there is a subset E^* of edges with $x(E^*) \geq 10^{-12}n$ such that for each $e \in E^*$ the probability that e is good is at least 10^{-24} ;*
- (b) *there are at least $\frac{19}{20}n$ edges e with $x_e \geq 1 - 10^{-7}$.*

The proof of this theorem is very long. It uses deeper results about random spanning trees and the structure of near-minimum cuts.

3.4 First improvement over Christofides for the GRAPHIC TSP

Following (56), we show now that Theorem 9 implies a better approximation ratio for the GRAPHIC TSP.

In case (a), Wolsey's analysis can be improved: let $y_e := x_e / (2 + 10^{-15})$ for good edges e and $y_e := x_e / 2$ for other edges e . Since $y(\delta(U)) \geq 1$ for every T_S -cut $\delta(U)$, we conclude that y is a feasible solution to (4). Therefore the expected cost of a minimum weight T_S -join is at most $c(y) \leq \frac{1}{2}c(x) - 10^{-16} \sum_{e \in E: e \text{ good}} c(e)x_e \leq \frac{1}{2}c(x) - 10^{-40} \sum_{e \in E^*} c(e)x_e$. If c is a graphic metric, we have $c(e) \geq 1$ for all $e \in E$ and $c(x) \leq 2n$. Then we get $c(y) \leq \frac{1}{2}c(x) - 10^{-40}x(E^*) \leq \frac{1}{2}c(x) - 10^{-52}n \leq \frac{1}{2}(1 - 10^{-52})c(x)$.

In case (b), the approximation ratio is better, and the proof is also easy. Let I be the set of edges e with $x_e \geq 1 - 10^{-7}$. The edges in I form vertex-disjoint paths and circuits, and each circuit has length

at least 10^7 (or is Hamiltonian). Remove one edge from each circuit and add edges of cost 1 to obtain a spanning tree (V, S) . Note that $c(S \setminus I) = |S \setminus I| < (\frac{1}{20} + 10^{-7})n \leq (\frac{1}{20} + 10^{-7})c(x)$. We get $c(S) = c(S \cap I) + c(S \setminus I) \leq \sum_{e \in S} c(e)x_e / (1 - 10^{-7}) + (\frac{1}{20} + 10^{-7})c(x)$.

Finally we add a minimum weight T_S -join J , where T_S is the set of vertices with odd degree in (V, S) . To bound $c(J)$, let $y_e := \frac{1}{3}$ for $e \in S$ and $y_e := \frac{2}{3}x_e$ for other edges e . We show that y is a feasible solution to (4).

For any set U with $|U \cap T_S|$ odd we have $|\delta(U) \cap S|$ odd. If $|\delta(U) \cap S| = 1$, then $y(\delta(U)) \geq \frac{1}{3} + y(\delta(U) \setminus S) \geq \frac{1}{3} + \frac{2}{3}(x(\delta(U)) - 1) \geq 1$. If $|\delta(U) \cap S| \geq 3$, then $y(\delta(U)) \geq 3 \cdot \frac{1}{3} = 1$.

Hence $c(J) \leq c(y) = \frac{1}{3}c(S) + \frac{2}{3} \sum_{e \in E \setminus S} c(e)x_e$. We conclude $c(S \cup J) \leq \frac{4}{3}c(S) + \frac{2}{3} \sum_{e \in E \setminus S} c(e)x_e \leq \frac{4}{3}c(x) / (1 - 10^{-7}) + \frac{4}{3}(\frac{1}{20} + 10^{-7})c(x) \leq (\frac{7}{5} + 10^{-6})c(x)$.

Note that we used properties of the GRAPHIC TSP in both cases, (a) and (b). Although the improvement over Christofides' algorithm is tiny (in case (a)), this result received a lot of interest.

4 Correcting parity by adding and removing edges

So far, all algorithms began with a spanning tree and then added edges to make the graph Eulerian. (53) had a brilliant idea: if we begin with a 2-connected graph, we may also delete some edges for making it Eulerian, and this may be cheaper overall.

4.1 Removable pairings

The following definition of (53) is very interesting. A *removable pairing* in a 2-vertex-connected graph (V, E) is a pair (R, \mathcal{P}) with the following properties:

- (a) $R \subseteq E$;
 - (b) for each $P \in \mathcal{P}$ there exists a vertex $v \in V$ and three distinct edges e_1, e_2, e_3 incident to v such that $P = \{e_1, e_2\}$;
 - (c) the elements of \mathcal{P} are pairwise disjoint;
 - (d) for any set $F \subseteq R$ with $|F \cap P| \leq 1$ for all $P \in \mathcal{P}$, the graph $(V, E \setminus F)$ is connected.
- (53) proposed to obtain a removable pairing as follows.

Lemma 10. Let $G = (V, E)$ be a 2-vertex-connected graph and (V, S) a DFS-tree in G , rooted at $r \in V$. For each edge $e = \{v, w\} \in E \setminus S$, let w.l.o.g. be v on the r - w -path in (V, S) , and let v' be the successor of v on this path. Add e to R ; moreover if $|\delta(v)| \geq 3$ and $e' = \{v, v'\}$ has not yet been added to R , then add also e' to R and $\{e, e'\}$ to \mathcal{P} (cf. Figure 3). Then (R, \mathcal{P}) is a removable pairing in G .

Proof. (a)–(c) are easy to see. To show that condition (d) holds, take $F \subseteq R$ with $|F \cap P| \leq 1$ for all $P \in \mathcal{P}$. For each $v \in V$ we consider the set W_v of vertices w for which v is on the r - w -path in (V, S) . We show that for each $v \in V$ the vertex set W_v induces a connected subgraph of $(V, E \setminus F)$. Indeed, this follows from a straightforward induction on $|W_v|$. \square




Figure 3. A 2-connected graph with a DFS tree (left, solid edges) and a removable pairing (right: dashed and dotted edges are in R ; arrows indicate pairs).




Figure 4. Proof of Theorem 11. The graph G' on the left results from G and (R, \mathcal{P}) in Figure 3. Squares denote odd-degree vertices. Here $|E| = 11$ and $|R| = 7$. As J' one could choose, e.g., the four edges whose weight is shown. This leads to the tour shown on the right.

4.2 The Mömke–Svensson lemma

Now we can formulate and prove the key lemma of (53). It works for general weights, although it has been used so far only for $c \equiv 1$. We follow the proof of (66), a variant of the original proof:

Theorem 11. Let $G = (V, E)$ be a 2-vertex-connected graph, $c : E(G) \rightarrow \mathbb{R}$, and (R, \mathcal{P}) a removable pairing in G . Then one can find a tour in G of length at most $\frac{4}{3}c(E) - \frac{2}{3}c(R)$ in $O(n^3)$ time.

Proof. Let T_G be the set of odd degree vertices of G . Let $c'(e) = c(e)$ for $e \in E \setminus R$ and $c'(e) = -c(e)$ for $e \in R$. For any T_G -join J in G that intersects each pair $P \in \mathcal{P}$ in at most one edge, we construct a tour from E by doubling the edges in $J \setminus R$ and deleting the edges in $J \cap R$. This tour has length $c(E) + c'(J)$.

To compute a T_G -join of weight at most $\frac{1}{3}c(E) - \frac{2}{3}c(R) = \frac{1}{3}c'(E)$, intersecting each pair at most once, we construct an auxiliary graph G' with weights c' from (G, c') as follows (cf. Figure 4). For each pair $P = \{\{v, w\}, \{v, w'\}\} \in \mathcal{P}$ we add a vertex v_P and an edge $\{v, v_P\}$ of weight zero, and replace the two edges in P by $\{v_P, w\}$ and $\{v_P, w'\}$, keeping their weight.

Let $T_{G'}$ be the set of odd degree vertices of G' . G' is 2-edge-connected. Hence every $T_{G'}$ -cut contains at least three edges, and the vector with all components $\frac{1}{3}$ is in the $T_{G'}$ -join polytope of G' (cf. Proposition 4), and even in its face defined by $x(\delta(v_P)) = 1$ for all $P \in \mathcal{P}$. Hence there is a $T_{G'}$ -join J' in G' with $|\delta_{J'}(v_P)| = 1$ for all $P \in \mathcal{P}$ and with weight at most $\frac{1}{3}c'(E)$. Such a J' can be found in $O(n^3)$ time (by adding a large weight to edges incident to v_P , for all $P \in \mathcal{P}$). It corresponds to a T_G -join J in G that intersects each pair at most once and has weight at most $\frac{1}{3}c'(E)$. \square

4.3 Subcubic graphs

Boyd, Sitters, van der Ster and Stougie [2011] devised a $\frac{4}{3}$ -approximation algorithm for cubic graphs. (53) gave a simpler proof for this result and extended it to subcubic graphs (i.e., graphs with maximum degree 3). Indeed, Lemma 10 yields a removable pairing with $|R| \geq 2(|E| - |S|) - 1$, because all non-tree edges, except possibly one incident to the root, can be paired with tree edges if the graph is subcubic. Theorem 11 yields a tour with at most $\frac{4}{3}|E| - \frac{2}{3}|R| = \frac{4}{3}n - \frac{2}{3}$ edges. This is best possible, e.g. for graphs that consist only of three internally vertex-disjoint paths of the same length and with the same endpoints.

(21) refined the techniques of (11); they can compute a tour of length less than $(\frac{4}{3} - \frac{1}{61236})n$ in any cubic graph in polynomial time.

4.4 Removable pairing via circulation

(53) showed how to find a good removable pairing in general graphs by a network flow approach, somewhat similar to Theorem 8. The idea again to start with a DFS tree and include some of the non-tree edges to make the subgraph 2-vertex-connected, but use as few as possible non-pairable edges.

First the input graph G is transformed into a flow network D as follows (cf. Figure 5). Let (V, S) be again a DFS tree, rooted at r . Note that r has degree 1 because G is 2-connected. Orient all tree




Figure 5. Flow network D in which we look for a circulation. Tree arcs (except the one incident to r) require at least one unit of flow. New vertices i_e ($e \in S$) are shown as squares.

edges away from r and all non-tree edges towards r . Subdivide each arc $e \in S$ by a vertex i_e . For each non-tree arc (v, w) add an arc (v, i_e) , where e is the first edge on the w - v -path in (V, S) .

Let $l((v, i_e)) := 1$ for each $v \in V \setminus \{r\}$ and $e \in \delta_S^+(v)$, and $l(e) := 0$ for all other arcs in D . Let $c(f) := \sum_{e \in S} \max\{0, f(\delta^-(i_e)) - 1\}$. (53) proved:

Lemma 12. *Let f be an integral circulation in D with $f \geq l$. Then one can construct a tour with at most $\frac{4}{3}n + \frac{2}{3}c(f) - \frac{2}{3}$ edges in $O(n^3)$ time.*

Proof. Let B be the set of non-tree edges that correspond to edges in D with positive flow. $(V, S \cup B)$ is 2-vertex-connected. Let $C := \{e = (v, w) \in S : v \neq r, f(\delta^-(i_e) \setminus \{(v, i_e)\}) > 0\}$ be the set of tree edges that can be paired (with a non-tree edge). Define a removable pairing in G by $R := B \cup C$ and letting \mathcal{P} contain a pair P for each element of C : for $e \in C$ choose an $e' \in B$ that corresponds to an edge in $\delta^-(i_e)$ and let $P = \{e, e'\}$. By Lemma 10, (R, \mathcal{P}) is indeed a removable pairing.

Now we apply Theorem 11 and obtain a tour with at most $\frac{4}{3}|S \cup B| - \frac{2}{3}|R| = \frac{4}{3}|S| + \frac{2}{3}|B| - \frac{2}{3}|C| = \frac{4}{3}(n - 1) + \frac{2}{3}c(f) + \frac{2}{3}$ edges. \square

An integral circulation in D with $f \geq l$ and $c(f)$ minimum can be found in $O(n^3)$ time. To bound the cost, (53) (and then also (55)) proceeded as follows.

1. Compute an optimum basic solution x of (6), with $c \equiv 1$ (in fact, (53) and (55) used (7) instead, but using (6) simplifies the proof; cf. Proposition 6).
2. Compute a DFS tree (V, S) by choosing a root arbitrarily and following always an edge e with maximum x_e to an unvisited vertex.
3. Define the following fractional circulation f' in the associated flow network (D, l) : For each $e \in E \setminus S$ send x_e units of flow along the fundamental cycle of e (the circuit in D corresponding to the unique circuit in $(V, S \cup \{e\})$).
4. For each $v \in V \setminus r$ and $e \in \delta^+(v)$ with $f'(v, i_e) < 1$, send $1 - f'(v, i_e)$ units of flow along any fundamental cycle containing e ; this circulation is called f'' . Let $f := f' + f''$. Then $f \geq l$.

(53) proved $c(f) \leq (4\sqrt{2} - 3)x(E) - (6\sqrt{2} - 6)n$. (55) improved the analysis and obtained $c(f) \leq \frac{5}{3}x(E) - \frac{3}{2}n$. The heart of his proof consists of showing that for each $v \in V \setminus \{r\}$ the contribution of the edges in $B_v := \{e \in \delta^-(v) : x_e > 0\}$ to $c(f)$ plus the extra flow added for v in step 4 is at most $\frac{1}{6}|B_v| + \frac{5}{6}(x(\delta(v)) - 2)$; the result then follows from summation, using the fact that x has at most $2n - 3$ nonzero variables (cf. Section 2.9).

We do not know whether Mucha's bound is tight. Together with Lemma 12 it directly yields a $\frac{13}{9}$ -approximation algorithm for the GRAPHIC TSP: we get a tour with at most $\frac{4}{3}n + \frac{10}{9}x(E) - n$ edges.

In the case that x in Step 1 is half-integral, we actually get a $\frac{4}{3}$ -approximation (as observed by (53)): we may assume that the support graph is 2-connected (otherwise consider its blocks separately); then $f'' \equiv 0$ and $c(f) = 0$. This is particularly interesting because (63) conjectured that the worst case for the integrality ratio occurs

when x is an optimum fractional perfect simple 2-matching (and hence w.l.o.g. half-integral).

5 Using ear-decompositions and matroids

5.1 Ear-decompositions

An ear-decomposition of a connected graph $G = (V, E)$ is a sequence P_0, P_1, \dots, P_k of subgraphs of G such that P_0 consists of a single vertex, $\{E(P_1), \dots, E(P_k)\}$ is a partition of E , and for $i = 1, \dots, k$, either P_i is a path with exactly its endpoints in $V(P_0) \cup \dots \cup V(P_{i-1})$ or P_i is a circuit with exactly one of its vertices (called its endpoint) in $V(P_0) \cup \dots \cup V(P_{i-1})$.

The vertices of an ear that are not endpoints are called its *internal vertices*. The length of an ear is the number of its edges; this is always the number of internal vertices plus one. An ear is called *trivial* if it has length 1, otherwise *nontrivial*. We call an ear *short* if it has length 2 or 3, otherwise *long*. An ear is called *odd* if its length is odd, otherwise even. The number of ears is always $|E| - |V| + 1$. See Figure 6, left-hand side, for an example.

(68) observed that a graph is 2-edge-connected if and only if it has an ear-decomposition. Hence computing an ear-decomposition with minimum number of nontrivial ears is equivalent to finding the smallest 2-edge-connected spanning subgraph (2ECSS); this problem is NP-hard. However, the number of even ears can be minimized in polynomial time. This is a fundamental result of (32) (also proved in Schrijver's [2003] book):

Theorem 13. *Let $G = (V, E)$ be a 2-edge-connected graph. Let $\varphi(G)$ denote the minimum number of even ears in any ear-decomposition of G . Then for any $T \subseteq V$ such that $|T|$ is even, there exists a T -join in G with at most $\frac{1}{2}(|V| + \varphi(G) - 1)$ edges. Moreover, there exists a $T \subseteq V$ such that $|T|$ is even and the minimum cardinality of a T -join in G is $\frac{1}{2}(|V| + \varphi(G) - 1)$. Such a T and an ear-decomposition with $\varphi(G)$ even ears can be found in $O(|V||E|)$ time.*

Any 2-edge-connected spanning subgraph (2ECSS) of G has at least $\varphi(G)$ ears in any ear-decomposition. Hence any 2ECSS, and thus any tour, has at least $n - 1 + \varphi(G)$ edges. (17) used Theorem 13 to strengthen this statement. Let

$$LP(G) := \min\{x(E) : x \geq 0, x(\delta(U)) \geq 2 \ (\emptyset \neq U \subset V)\}. \quad (9)$$

Note that (9) is the special case of (7) for $c(e) = 1$ for all $e \in E$, and $LP(G)$ is a lower bound on the length of an optimum tour.

Corollary 14. *For any 2-edge-connected graph G we have*

$$L_\varphi := n - 1 + \varphi(G) \leq LP(G).$$

Proof. By Theorem 13 there exists a set T of vertices such that $|T|$ is even and $\frac{1}{2}(n - 1 + \varphi(G))$ is the minimum cardinality of a T -join in G . Now observe that (4) is at most half of (7). \square




Figure 6. Left: A graph G with an ear-decomposition. The internal vertices of the i -th ear are labelled i . We have $\varphi(G) = 2$; ears 1 and 5 are even. Ears 3, 4, 5, and 6 are short; they are all pendant. Dotted edges are trivial ears. Right: Trivial ears are deleted, and a removable pairing (R, \mathcal{P}) with $|R| = 8$ and $|\mathcal{P}| = 2$ as in the proof of Theorem 15 is shown.

5.2 Applying the Mömke-Svensson lemma to ear-decompositions

Let us call an ear *pendant* if none of its internal vertices is endpoint of any nontrivial ear. By applying Theorem 11 to an ear-decomposition of a graph G , (66) observed:

Theorem 15. Given a 2-vertex-connected graph G with an ear-decomposition with π pendant ears and no trivial ears, one can construct a tour with at most $\frac{4}{3}(n-1) + \frac{2}{3}\pi$ edges in $O(n^3)$ time.

Proof. Define a removable pairing by taking an arbitrary edge of each pendant ear, and for each other ear a pair of its edges, incident to a common vertex that is endpoint of another ear. If $k = |E| - |V| + 1$ denotes the number of ears, we have $|R| = 2k - \pi$. Theorem 11 yields a tour with at most $\frac{4}{3}|E| - \frac{2}{3}|R| = \frac{4}{3}(n-1) + \frac{2}{3}\pi$ edges. \square

See Figure 6 for an example. This bound is good if there are few pendant ears. Otherwise we need something else. It turns out that long pendant ears are easy to deal with, but short pendant ears require care.

5.3 Nice and nicer ear-decompositions

Frank's Theorem 13 was also used by (17) and (66) as a starting point to obtain a nice ear-decomposition. An ear-decomposition is called *nice* if it has $\varphi(G)$ even ears, all short ears are pendant, and there is no edge joining internal vertices of different short ears. (Figure 6, left-hand side, displays a nice ear-decomposition.)

Lemma 16. Given a 2-vertex-connected graph, one can compute a nice ear-decomposition in polynomial time.

A nice ear-decomposition allows for optimizing the short ears in the following sense. Let M contain for each short ear the set of its internal vertices (cf. Figure 7, left-hand side). For $f \in M$ we denote by E_f the set of pairs $\{v, w\}$ such that G contains a path from v to w whose set of internal vertices is f . We will pick an $e_f \in E_f$ for each $f \in M$ such that $(V, \{e_f : f \in M\})$ has as few connected components as possible.

Equivalently, we pick an $e_f \in E_f$ for each $f \in M$ such that the rank of $\{e_f : f \in M\}$ in the graphic matroid is maximum. Denote this maximum by μ . By Rado's [1942] Theorem,

$$\mu = \min\{r(\cup_{i \in I} E_i) + |M \setminus I| : I \subseteq M\}. \tag{10}$$

(In the example of Figure 7, $|M| = 4, \mu = 3$ and $I = \{5, 6\}$ attains the minimum.) The maximum can be found by an algorithm for matroid intersection. (66) found an $O(|V||E|)$ -time algorithm.

We then replace the short ear with internal vertices f by a path with internal vertices f and endpoints e_f , for each $f \in M$. This may change the set of trivial ears, but the ear-decomposition remains nice. Let V_M denote the union of the sets in M (i.e., the set of internal vertices of short ears). We have:

Theorem 17. Let G be a 2-connected graph with a nice ear-decomposition. Then we can compute in polynomial time another nice ear-decomposition of G such that (V, F) , where F contains all edges of short ears, has $|V| - |V_M| - \mu$ connected components. \square




Figure 7. Left: Optimizing the ear-decomposition of Figure 6. The elements of M are the gray sets. Here only ear 4 needed to be replaced. Right: The cuts in the proof of Theorem 18.

We also get another lower bound:

Theorem 18. Let G be a 2-connected graph with a nice ear-decomposition with $|M|$ short ears. Then

$$L_\mu := n - 1 + |M| - \mu \leq \text{LP}(G).$$

Proof. Let I be a set attaining the minimum in (10). Let \mathcal{U} be the partition of V such that for each $f \in I$ there is a $U \in \mathcal{U}$ with $f \subseteq U$ and $e_f \subseteq U$ for all $e_f \in E_f$, and $|\mathcal{U}| = |V| - |V_I| - r(\cup_{i \in I} E_i)$. By Rado's Theorem (cf. (10)) we have $|\mathcal{U}| = |V| - |V_I| - \mu + |M \setminus I|$.

Consider the family of sets $\mathcal{U} \cup I \cup \{\{v\} : v \in V_I\}$, taking singletons in I twice. Summing over the inequalities $x(\delta(U)) \geq 2$ for these sets U (unless $U = V$) completes the proof because no edge is contained in more than two of these at least $n - 1 + |M| - \mu$ cuts. \square

See Figure 7 (right-hand side) for an illustration. The set F in Theorem 17 consists of the black edges in Figure 8, left-hand side.

5.4 The $\frac{7}{5}$ -approximation algorithm

Now we can explain the $\frac{7}{5}$ -approximation algorithm for the GRAPHIC TSP by (66). Let $\Lambda^G := \frac{2}{3}L_\mu + \frac{1}{3}L_\varphi$. Note that Λ^G is a lower bound on the optimum and in fact on $\text{LP}(G)$, and $\Lambda^G \geq n - 1$ (cf. Corollary 14 and Theorem 18). We first show:

Lemma 19. Let G be a 2-vertex-connected graph with a nice ear-decomposition that has no trivial ears and for which the union of all short ears have $|V| - |V_M| - \mu$ connected components. Then one can compute a tour in G with at most $\frac{7}{5}\Lambda$ edges in $O(n^3)$ time.

Proof. If $\pi \leq \frac{\Lambda}{10}$, apply Theorem 15 and obtain a tour with at most $\frac{4}{3}(n-1) + \frac{2}{3}\pi \leq \frac{7}{5}\Lambda$ edges.

Otherwise take all $|V_\pi| + \pi$ edges of pendant ears, where V_π denotes the internal vertices of pendant ears. Add at most $n - |V_\pi| - \mu - 1$ edges of $G[V \setminus V_\pi]$ to obtain a connected spanning subgraph. Let $T \subseteq V \setminus V_\pi$ be the set of vertices with odd degree in this subgraph. Then add a minimum T -join in $G[V \setminus V_\pi]$; by Theorem 13 it has at most $\frac{1}{2}(n - |V_\pi| - 1 + \varphi - \varphi_\pi)$ edges. Summing up, our tour has at most $\pi + n - \mu - 1 + \frac{1}{2}(n - |V_\pi| - 1 + \varphi - \varphi_\pi)$ edges. Observing $|V_\pi| \geq 4\pi - 2|M| - \varphi_\pi$, this is at most $L_\mu + \frac{1}{2}L_\varphi - \pi \leq \frac{7}{5}\Lambda$. \square

Figure 8 illustrates the second part of this proof for the ear-decomposition in Figure 7 (left-hand side).

The overall algorithm is now easily described:

1. Compute a nice ear-decomposition of G (Lemma 16).
2. Optimize the short ears (Theorem 17).
3. Delete all trivial ears.
4. Apply Lemma 19 to each block of the remaining graph.

It is not difficult to show that the sum of the lower bounds Λ for all blocks equals the lower bound Λ for G . This implies that we have a $\frac{7}{5}$ -approximation algorithm.




Figure 8. Left: The edges of pendant ears after optimizing short ears (black) and a minimal set of edges of non-pendant ears (gray) to make a connected spanning subgraph. Right: For correcting parities we need three more edges; a possible resulting tour is shown.

6 The path version and connected T -joins

What if we do not require the walk to be closed? Then we look for a (Hamiltonian) path in the metric closure. We assume that the endpoints are given (otherwise we can try all pairs, or take a tour and delete one edge) and distinct: the path must begin in s and end in t (where $s, t \in V$ and $s \neq t$). For any of the problems studied in this paper, this variant is called the s - t -path version or simply the path version.

6.1 Asymmetric path version

Obviously, any ρ -approximation for the s - t -path version implies a ρ -approximation algorithm for the ASYMMETRIC TSP itself: just guess any edge (t, s) in an optimum solution (fix s and try all $n-1$ possibilities for t). (29) showed that the opposite also holds approximately: any ρ -approximation for the ASYMMETRIC TSP implies a $(2 + \epsilon)\rho$ -approximation algorithm for the path version.

6.2 Undirected path version and connected T -joins

In the undirected case, if we ask for a walk from s to t , by Section 1.3 this is equivalent to ask for a connected spanning multi-subgraph in which s and t have odd degree and all other vertices have even degree. It is natural to generalize this further to prescribe arbitrary parities: the CONNECTED T -JOIN PROBLEM asks for a set F such that (V, F) is a connected graph and F is a T -join. We call such a set F simply a *connected T -join*, or a *T -tour*. Again, F may contain pairs of parallel edges.

For $T = \emptyset$ this is the SYMMETRIC TSP, and for $T = \{s, t\}$ this is its s - t -path version. Again we may consider the GRAPHIC special case, where all edges of F must be copies of edges of the input graph.

Christofides' [1976] algorithm also works for the CONNECTED T -JOIN PROBLEM: take a minimum spanning tree (V, S) and add a minimum-weight T_S -join, where T_S is now the set of vertices whose degree in (V, S) has the wrong parity (so S is a $(T_S \Delta T)$ -join).

However, this generalization of Christofides' algorithm is only a $\frac{5}{3}$ -approximation algorithm ((43; 66)). To see this, let (V, S) be a minimum spanning tree. Let J be a minimum weight T_S -join, and J^* an optimum solution (a minimum weight connected T -join). Then $S \dot{\cup} J^*$ is a T_S -join. Since both S and J^* are connected, each contains a T_S -join; so $S \dot{\cup} J^*$ can be partitioned into three T_S -joins. Hence $3c(S \dot{\cup} J) \leq 3c(S) + 3c(J) \leq 3c(S) + c(S \dot{\cup} J^*) = 4c(S) + c(J^*) \leq 5c(J^*)$. The bound is tight even for $|T| = 2$ ((43)), as the graphs $(\{0, \dots, 3k\}, \{\{i, i+1\} : 0 \leq i < 3k\} \cup \{\{3i, 3i+3\} : 0 \leq i < k\})$ show.

(66) showed that their techniques (outlined in Section 5 above) also lead to a $\frac{3}{2}$ -approximation algorithm for the GRAPHIC CONNECTED T -JOIN PROBLEM. Previously, there were only algorithms for the special case $|T| = 2$; here the best was the 1.578-approximation algorithm of (2).

6.3 Best-of-many Christofides' algorithm

(2) also found a 1.619-approximation algorithm for the path version ($|T| = 2$) with general weights. This was the first improvement of Christofides' algorithm that is not restricted to the graphic special case. This algorithm was generalized by (16); they obtain an approximation ratio of 1.625 for $|T| \geq 4$. Then (65) obtained an $\frac{8}{5}$ -approximation algorithm for arbitrary T and general weights.

All these three papers analyze essentially the same algorithm, which (2) called *best-of-many Christofides*: it computes an optimum solution to a natural LP relaxation (see (11) below) and writes it as convex combination of spanning trees (plus a nonnegative vector). For each of these spanning trees, S , we again compute a minimum weight T_S -join J , where T_S is the set of vertices of S whose degree

has the wrong parity, and output the best of these T -tours $S \dot{\cup} J$.

Following (66) and (65), we consider the LP relaxation

$$\begin{aligned} \min \quad & c(x) \\ \text{subject to} \quad & x(\delta(U)) \geq 2 \quad (\emptyset \neq U \subset V, |U \cap T| \text{ even}) \\ & x(\delta(W)) \geq |W| - 1 \quad (W \text{ partition of } V) \\ & x_e \geq 0 \quad (e \in E) \end{aligned} \tag{11}$$

Here $\delta(W)$ denotes the set of edges with endpoints in different classes of the partition W . For an optimum solution x (in fact for every feasible solution) we can write $x \geq \sum_{S \in \mathcal{S}} p_S \chi^S$, where again S denotes the set of edge sets of spanning trees, χ^S denotes the incidence vector of S , and $p_S \geq 0$ for all $S \in \mathcal{S}$ and $\sum_{S \in \mathcal{S}} p_S = 1$. ((2) and (16) work in the metric closure and use a stronger LP in order to obtain $x = \sum_{S \in \mathcal{S}} p_S \chi^S$, but this is not necessary.)

By Carathéodory's theorem we can assume that $p_S > 0$ for less than n^2 spanning trees (V, S) . An optimum LP solution x , such spanning trees, and such numbers p_S can be computed in polynomial time, as can be shown with the ellipsoid method. For each $S \in \mathcal{S}$ with $p_S > 0$, the algorithm computes a minimum weight T_S -join J and consider the T -tour $S \dot{\cup} J$; we output the best of these. Its cost is $\min_{S \in \mathcal{S}: p_S > 0} (c(S) + \min\{c(J) : J \text{ is a } T_S\text{-join}\}) \leq \sum_{S \in \mathcal{S}} p_S (c(S) + \min\{c(J) : J \text{ is a } T_S\text{-join}\}) \leq c(x) + \sum_{S \in \mathcal{S}} p_S c(\chi^S)$, for any set of vectors $(\chi^S)_{S \in \mathcal{S}}$ such that χ^S is in the T_S -join polyhedron (cf. (4)). The difficulty in the analysis lies in finding an appropriate set of vectors $(\chi^S)_{S \in \mathcal{S}}$.

6.4 Analysis

Let $\mathcal{Q} := \{Q = \delta(U) : \emptyset \neq U \subset V, x(Q) < 2\}$. (2) proposed to choose $\chi^S := (1 - 2\beta)\chi^S + \beta x + r^S$, where $\beta \leq \frac{1}{2}$ and r^S is a nonnegative vector satisfying

$$r^S(Q) \geq 4\beta - 1 - \beta x(Q) \tag{12}$$

for all $S \in \mathcal{S}$ and all $Q \in \mathcal{Q}$ with $|Q \cap S| \geq 2$.

Then for each $S \in \mathcal{S}$ and each T_S -cut Q we have $\chi^S(Q) \geq 1$. Indeed, if $Q \notin \mathcal{Q}$, then $\chi^S(Q) \geq (1 - 2\beta)|S \cap Q| + \beta x(Q) \geq 1 - 2\beta + 2\beta = 1$. If $Q \in \mathcal{Q}$, then Q is not only a T_S -cut but also a T -cut, so $|Q \cap S|$ is even and hence at least two, and we have $\chi^S(Q) = (1 - 2\beta)|Q \cap S| + \beta x(Q) + r^S(Q) \geq 2 - 4\beta + \beta x(Q) + r^S(Q) \geq 1$.

So χ^S is in the T_S -join polyhedron for all $S \in \mathcal{S}$. Moreover, $\sum_{S \in \mathcal{S}} p_S c(\chi^S) \leq (1 - 2\beta) \sum_{S \in \mathcal{S}} p_S c(S) + \beta c(x) + \sum_{S \in \mathcal{S}} p_S c(r^S) \leq (1 - \beta)c(x) + \sum_{S \in \mathcal{S}} p_S c(r^S)$.

(2) chose $\beta = 1/\sqrt{5}$ and found a vector r such that $r^S = r$ for all $S \in \mathcal{S}$ satisfies (12) and $c(r) \leq (7\sqrt{5} - 15)/10$, yielding the approximation ratio $(1 + \sqrt{5})/2$ (the golden ratio).

(65) improved this by letting $v^Q := \sum_{S \in \mathcal{S}: |Q \cap S|=1} p_S \chi^{Q \cap S}$ for $Q \in \mathcal{Q}$, and

$$r^S := \sum_{Q \in \mathcal{Q}: |Q \cap S| \geq 2} \max \left\{ 0, \frac{4\beta - 1 - \beta x(Q)}{2 - x(Q)} \right\} v^Q.$$

Note that $v^Q(Q) = \sum_{S \in \mathcal{S}: |Q \cap S|=1} p_S \geq 2 - \sum_{S \in \mathcal{S}} p_S |Q \cap S| \geq 2 - x(Q)$. To show (12), simply observe that for $S \in \mathcal{S}$ and $Q \in \mathcal{Q}$ with $|Q \cap S| = 2$ we have $r^S(Q) \geq \frac{4\beta - 1 - \beta x(Q)}{2 - x(Q)} v^Q(Q)$.

We now bound the cost. Note that $\sum_{S \in \mathcal{S}: |Q \cap S| \geq 2} p_S \leq x(Q) - 1$ for $Q \in \mathcal{Q}$. Using this, we obtain the bound $\sum_{S \in \mathcal{S}} p_S c(r^S) \leq \sum_{Q \in \mathcal{Q}} (x(Q) - 1) \max \left\{ 0, \frac{4\beta - 1 - \beta x(Q)}{2 - x(Q)} \right\} c(v^Q) \leq \sum_{Q \in \mathcal{Q}} \frac{1}{9} c(v^Q) = \frac{1}{9} \sum_{S \in \mathcal{S}} p_S \sum_{Q \in \mathcal{Q}: |Q \cap S|=1} c(Q \cap S) \leq \frac{1}{9} \sum_{S \in \mathcal{S}} p_S c(S \setminus J_S)$, where J_S denotes the unique subset of S that is a T_S -join.

Let $\beta = \frac{4}{9}$. If $\sum_{S \in \mathcal{S}} p_S c(S \setminus J_S) \leq \frac{2}{5} c(x)$, we have $\sum_{S \in \mathcal{S}} p_S c(\chi^S) \leq (1 - \frac{4}{9} + \frac{2}{45}) c(x) = \frac{3}{5} c(x)$. Otherwise $\sum_{S \in \mathcal{S}} p_S c(J_S) \leq \frac{3}{5} c(x)$, and noting that χ^{J_S} is in the T_S -join polyhedron, we get the same bound. This gives Sebő's [2012] $\frac{8}{5}$ -approximation algorithm.

7 Further results

We briefly mention some other related results. However, it is impossible to mention all important results here.

7.1 Inapproximability

(50) proved that no $\frac{185}{184}$ -approximation algorithm exists for the SYMMETRIC TSP unless $P = NP$. (57) proved that no $\frac{118}{117}$ -approximation algorithm exists for the ASYMMETRIC TSP unless $P = NP$.

7.2 Integrality Ratios

We have seen in Section 2.2 that the integrality ratio of (2) is at least $\frac{4}{3}$ even for graphic metrics. The integrality ratio of (2) is conjectured to be exactly $\frac{4}{3}$ even for general metrics, but this so-called *TSP- $\frac{4}{3}$ -conjecture* is open; we only know Wolsey's [1980] upper bound $\frac{3}{2}$ in general (cf. Section 2.5) and the upper bound $\frac{7}{5}$ for graphic metrics by (66) (cf. Section 5).

The *TSP- $\frac{4}{3}$ -conjecture* is supported by computational verification for $n \leq 12$ ((10)) and by theoretical work of (37), who proved that for any instance with ratio greater than $\frac{4}{3}$ even the LP that arises from (7) by adding many classes of inequalities that are valid for the graphical traveling salesman polyhedron does not have an integral optimum solution.

(63) showed that the worst ratio of an optimum perfect 2-matching over (2) is $\frac{10}{9}$, as conjectured by (9).

For (6) and (7), the integrality ratio is between $\frac{6}{5}$ and $\frac{3}{2}$ ((1)). (12) conjectured it to be $\frac{4}{3}$. The ratio restricted to unit weights is between $\frac{9}{8}$ and $\frac{4}{3}$ ((66)).

The integrality ratio of the asymmetric subtour LP is at least 2 (shown by (14), disproving a conjecture of (13)) and at most $2 + 8 \ln n / \ln \ln n$ ((5)). The same holds for the path version ((34)).

7.3 Further special cases

An even more special case than the GRAPHIC TSP is the 1-2-TSP, in which $c(v, w) \in \{1, 2\}$ for all $v, w \in V$. To see that this is essentially a special case of the GRAPHIC TSP (up to an additive constant of 1), add a vertex x and consider the graph $(V \cup \{x\}, \{\{v, x\} : v \in V\} \cup \{\{v, w\} : v, w \in V, u \neq v, c(v, w) = 1\})$. The 1-2-TSP has an $\frac{8}{7}$ -approximation algorithm ((7)) but no $\frac{744}{743}$ -approximation algorithm ((27)). The integrality ratio of (2) for the 1-2-TSP is between $\frac{10}{9}$ and $\frac{19}{15}$ ((59)).

In the special case of the GRAPHIC TSP where the instance is a k -regular graph (with k large), (67) showed how to find a tour of length at most $(1 + \sqrt{64/\ln k})n$ in polynomial time.

7.4 Geometric instances and planar graphs

(4) found an approximation scheme for geometric instances. Here, each city is associated with a point in \mathbb{R}^d , and the distances are ℓ_p -distances. This case is also NP-hard, for any fixed $d \geq 2$ and any p . The most prominent case $d = p = 2$ is called the EUCLIDEAN TSP (see also (52)). (61) improved the running time: for every fixed $\epsilon > 0$ they have a $(1 + \epsilon)$ -approximation algorithm that runs in $O(n \log n)$ time. However, the constants involved are still quite large for reasonable values of ϵ , and thus the practical value seems to be limited. (6) found a randomized approximation scheme for metric spaces with bounded doubling dimension.

For planar graphs with nonnegative edge weights, (48) found an approximation scheme that has linear running time for every fixed $\epsilon > 0$. An approximation scheme exists even for bounded genus graphs ((23)).

Interestingly, it is not known whether the decision version of the EUCLIDEAN TSP belongs to NP.

7.5 Polyhedral Descriptions

Many classes of facets of the TSP polytope have been discovered, but a complete description is out of reach. Recently, (31) proved that every polyhedron that projects to the TSP polytope (i.e., any extended formulation) has $2^{\Omega(n^{1/4})}$ facets. It may not be surprising that the TSP has no compact extended formulation, but this was not known before, and this result is unconditional (i.e., it does not assume $P \neq NP$). The proof reveals an interesting connection to communication complexity.

7.6 The 2ECSS problem

The integral solutions to (6) are the 2-edge-connected spanning subgraphs (2ECSS). (54) showed that the smallest 2ECSS can be smaller than the shortest tour by up to a factor $\frac{4}{3}$; this also follows directly from applying Theorem 11 to each block of a smallest 2ECSS. The bound is tight as Figure 2 shows.

(66) observed that the techniques of Section 5 directly imply a $\frac{4}{3}$ -approximation algorithm for the (unweighted) 2ECSS problem. Indeed, if $\pi \geq \frac{1}{6} \text{LP}(G)$, the second case of the proof of Lemma 19 yields a tour (and hence a 2ECSS) of length $\frac{4}{3} \text{LP}(G)$. Otherwise one can simply take all k nontrivial ears: we get $n - 1 + k$ edges, and this is at most $\frac{5}{4} L_\varphi + \frac{\pi}{2}$ since $n - 1 \geq 4k - 2\pi - \varphi(G)$.

Better approximation ratios have been claimed, but no complete proof has been published. (30) proved that the problem is MAXSNP-hard. For the weighted case, (47) found a 2-approximation algorithm, which is still the best known.

(66) also showed the following: if there is a ρ -approximation algorithm for the unweighted 2ECSS problem, then there is a $\frac{2}{3}(\rho + 1)$ -approximation algorithm for the GRAPHIC TSP.

8 Open problems

We conclude this survey by listing some open research problems that we consider important. Almost all of these problems have been formulated earlier, and indeed most of them are very natural. None of them seems to be easy. However, given the remarkable progress that has been made during the last few years, one may hope that we will see some solutions soon.

1. Improve Christofides' algorithm: find a ρ -approximation algorithm for the SYMMETRIC TSP for some $\rho < \frac{3}{2}$.
2. Find a constant-factor approximation algorithm for the ASYMMETRIC TSP, or at least the special case in which a strongly connected digraph (V, E) is given and $c(v, w) = 1$ if $(v, w) \in E$ and $c(v, w) = \infty$ otherwise (one might call this the DIGRAPHIC TSP).
3. Determine the integrality ratio of the subtour relaxation (2) of the SYMMETRIC TSP.
4. Prove a better bound on the integrality ratio for another (polynomial-time solvable) LP relaxation of the SYMMETRIC TSP.
5. Solve the LPs (2), (6), (7), and (8) by combinatorial algorithms.
6. Answer the question whether the integrality ratio of the directed subtour relaxation (8) is bounded by a constant.
7. How good is the best-of-many Christofides' algorithm (cf. Section 6.3) really; i.e., what is the worst case? The answer can of course be different for $T = \emptyset$ (the SYMMETRIC TSP) and for general T .
8. Improve the lower bounds on the approximability substantially.
9. Find a $\frac{4}{3}$ -approximation algorithm for the GRAPHIC TSP.
10. Find a $\frac{3}{2}$ -approximation algorithm for the CONNECTED T -JOIN PROBLEM with arbitrary nonnegative weights, at least in the special case $|T| = 2$.

11. Improve on the 2-approximation algorithm for the weighted 2ECSS problem. Note that Wolsey's analysis (Section 2.5) shows that Christofides' algorithm is also a $\frac{3}{2}$ -approximation algorithm for the variant of the 2ECSS problem where doubling edges is allowed. However, in contrast to the unweighted special case, allowing to double edges really changes the problem.

Acknowledgement. Thanks to Corinna Gottschalk, Swati Gupta, Satoru Iwata, Volker Kaibel, Marcin Mucha, R. Ravi, András Sebő, David Shmoys, and Ola Svensson for careful reading and useful remarks.

Jens Vygen, Research Institute for Discrete Mathematics, University of Bonn, Germany. www.or.uni-bonn.de/~vygen

References

- [1] Alexander, A., Boyd, S., and Elliott-Magwood, P. [2006]: On the integrality gap of the 2-edge connected subgraph problem. Technical Report TR-2006-04, SITE, University of Ottawa 2006
- [2] An, H.-C., Kleinberg, R., and Shmoys, D.B. [2012]: Improving Christofides' algorithm for the s-t path TSP. Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012), 875–886
- [3] Applegate, D.L., Bixby, R., Chvátal, V., and Cook, W.J. [2006]: The Traveling Salesman Problem: A Computational Study. Princeton University Press 2006
- [4] Arora, S. [1998]: Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. Journal of the ACM 45 (1998), 753–782
- [5] Asadpour, A., Goemans, M.X., Mądry, A., Oveis Gharan, S., and Saberi, A. [2010]: An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010), 379–389
- [6] Bartal, Y., Gottlieb, L.-A., and Krauthgamer, R. [2012]: The traveling salesman problem: low-dimensionality implies a polynomial time approximation scheme. Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012), 663–672
- [7] Berman, P., and Karpinski, M. [2006]: 8/7-approximation algorithm for (1,2)-TSP. Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2006), 641–648
- [8] Bläser, M. [2003]: A new approximation algorithm for the asymmetric TSP with triangle inequality. Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2003), 638–645
- [9] Boyd, S., and Carr, R. [2011]: Finding low cost TSP and 2-matching solutions using certain half-integer subtour vertices. Discrete Optimization 8 (2011), 525–539
- [10] Boyd, S., and Elliott-Magwood, P. [2007]: Structure of the extreme points of the subtour elimination polytope of the STSP. Technical Report TR-2007-09, SITE, University of Ottawa, Ottawa, 2007
- [11] Boyd, S., Sitters, R., van der Ster, S., and Stougie, L. [2011]: TSP on cubic and subcubic graphs. In: Integer Programming and Combinatorial Optimization; Proceedings of the 15th International IPCO Conference; LNCS 6655 (O. Günlük, G.J. Woeginger, eds.), Springer, Berlin 2011, pp. 65–77
- [12] Carr, R., and Ravi, R. [1998]: A new bound for the 2-edge connected subgraph problem. In: Integer Programming and Combinatorial Optimization; Proceedings of the 6th International IPCO Conference; LNCS 1412 (R.E. Bixby, E.A. Boyd, R.Z. Rios-Mercado, eds.), Springer, Berlin 1998; pp. 112–125
- [13] Carr, R., and Vempala, S. [2004]: On the Held-Karp relaxation for the asymmetric and symmetric traveling salesman problems. Mathematical Programming A 100 (2004), 569–587
- [14] Charikar, M., Goemans, M.X., and Karloff, H. [2006]: On the integrality ratio for the asymmetric traveling salesman problem. Mathematics of Operations Research 31 (2006), 245–252
- [15] Chekuri, C., Vondrák, J., and Zenklus, R. [2010]: Dependent randomized rounding via exchange properties of combinatorial structures. Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2010), 575–584
- [16] Cheriyan, J., Friggstad, Z., and Gao, Z. [2012]: Approximating minimum-cost connected T -joins. Approximation, Randomization, and Combinatorial Optimization; Proceedings of the 16th APPROX Workshop; LNCS 7408 (A. Gupta, K. Jansen, J. Rolim, R. Servedio, eds.), Springer, Berlin 2012, pp. 110–121
- [17] Cheriyan, J., Sebő, A., and Szigeti, Z. [2001]: Improving on the 1.5-approximation of a smallest 2-edge connected spanning subgraph. SIAM Journal on Discrete Mathematics 14 (2001), 170–180
- [18] Christofides, N. [1976]: Worst-case analysis of a new heuristic for the traveling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie-Mellon University, Pittsburgh 1976
- [19] Cook, W.J. [2012]: In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation. Princeton University Press 2012
- [20] Cornuéjols, G., Fonlupt, J., and Naddef, D. [1985]: The traveling salesman problem on a graph and some related integer polyhedra. Mathematical Programming 33 (1985), 1–27
- [21] Correa, J.R., Larré, O., and Soto, J.A. [2012]: TSP tours in cubic graphs: beyond 4/3. In: Algorithms – ESA 2012; LNCS 7501 (L. Epstein, P. Ferragina, eds.), Springer, Berlin 2012, pp. 790–801
- [22] Dantzig, G.B., Fulkerson, D.R., and Johnson, S.M. [1954]: Solution of a large scale traveling salesman problem. Operations Research 2 (1954), 393–410
- [23] Demaine, E.D., Hajiaghayi, M., and Mohar, B. [2010]: Approximation algorithms via contraction decomposition. Combinatorica 30 (2010), 533–552
- [24] Edmonds, J. [1965]: The Chinese postman's problem. Bulletin of the Operations Research Society of America 13 (1965), B-73
- [25] Edmonds, J. [1970]: Submodular functions, matroids and certain polyhedra. In: Combinatorial Structures and Their Applications; Proceedings of the Calgary International Conference on Combinatorial Structures and Their Applications 1969 (R. Guy, H. Hanani, N. Sauer, J. Schönheim, eds.), Gordon and Breach, New York 1970, pp. 69–87
- [26] Edmonds, J., and Johnson, E.L. [1973]: Matching, Euler tours and the Chinese postman. Mathematical Programming 5 (1973), 88–124
- [27] Engebretsen, L., and Karpinski, M. [2006]: TSP with bounded metrics. Journal of Computer and System Sciences 72 (2006), 509–546
- [28] Euler, L. [1736]: Solutio problematis ad geometriam situs pertinentis. Commentarii Academiae Petropolitanae 8 (1736), 128–140
- [29] Feige, U., and Singh, M. [2007]: Improved approximation algorithms for traveling salesperson tours and paths in directed graphs. Proceedings of the 10th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems; LNCS 4627 (M. Charikar, K. Jansen, O. Reingold, J.D.P. Rolim, eds.), Springer, Berlin 2007, pp. 104–118
- [30] Fernandes, C.G. [1998]: A better approximation ratio for the minimum size k -edge-connected spanning subgraph problem. Journal of Algorithms 28 (1998), 105–124
- [31] Fiorini, S., Massar, S., Pokutta, S., Tiwary, H.R., and de Wolf, R. [2012]: Linear vs. semidefinite extended formulations: exponential separation and strong lower bounds. Proceedings of the 44th Annual ACM Symposium on Theory of Computing (STOC 2012), 95–106
- [32] Frank, A. [1993]: Conservative weightings and ear-decompositions of graphs. Combinatorica 13 (1993), 65–81
- [33] Frieze, A.M., Galbiati, G., and Maffioli, F. [1982]: On the worst-case performance of some algorithms for the asymmetric traveling salesman problem. Networks 12 (1982), 23–39
- [34] Friggstad, Z., Gupta, A., and Singh, M. [2012]: An improved integrality gap for asymmetric TSP paths. Accepted for IPCO 2013
- [35] Gamarnik, D., Lewenstein, M., and Sviridenko, M. [2005]: An improved upper bound for the TSP in cubic 3-edge-connected graphs. Operations Research Letters, 33 (2005), 467–474
- [36] Garey, M.R., Johnson, D.S., and Tarjan, R.E. [1976]: The planar Hamiltonian circuit problem is NP-complete. SIAM Journal on Computing 5 (1976), 704–714
- [37] Goemans, M.X. [1995]: Worst-case comparison of valid inequalities for the TSP. Mathematical Programming, 69 (1995), 335–349
- [38] Goemans, M.X. [2006]: Minimum bounded-degree spanning trees. Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006), 273–282
- [39] Goemans, M.X., and Bertsimas, D.J. [1993]: Survivable networks, linear programming relaxations and the parsimonious property. Mathematical Programming 60 (1993), 145–166
- [40] Held, M., and Karp, R.M. [1970]: The traveling-salesman problem and minimum spanning trees. Operations Research 18 (1970), 1138–1162
- [41] Hierholzer, C. [1873]: Über die Möglichkeit, einen Linienzug ohne Wiederholung und ohne Unterbrechung zu umfahren. Mathematische Annalen 6 (1873), 30–32
- [42] Hoffman, A.J. [1960]: Some recent applications of the theory of linear inequalities to extremal combinatorial analysis. In: Combinatorial Analysis (R.E. Bellman, M. Hall, eds.), AMS, Providence 1960, pp. 113–128
- [43] Hoogeveen, J.A. [1991]: Analysis of Christofides' heuristic: some paths are more difficult than cycles. Operations Research Letters 10 (1991), 291–295

- [44] Kaplan, H., Lewenstein, M., Shafir, N., and Sviridenko, M. [2005]: Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. *Journal of the ACM* 52 (2005), 602–626
- [45] Karger, D.R., and Stein, C. [1996]: A new approach to the minimum cut problem. *Journal of the ACM* 43 (1996), 601–640
- [46] Karp, R.M. [1972]: Reducibility among combinatorial problems. In: *Complexity of Computer Computations* (R.E. Miller, J.W. Thatcher, eds.), Plenum Press, New York 1972, pp. 85–103
- [47] Khuller, S., and Vishkin, U. [1994]: Biconnectivity approximations and graph carvings. *Journal of the ACM* 41 (1994), 214–235
- [48] Klein, P.N. [2008]: A linear-time approximation scheme for TSP in undirected planar graphs with edge-weights. *SIAM Journal on Computing* 37 (2008), 1926–1952
- [49] Korte, B., and Vygen, J. [2012]: *Combinatorial Optimization: Theory and Algorithms*. Springer, Berlin, Fifth Edition 2012
- [50] Lampis, M. [2012]: Improved inapproximability for TSP. *Approximation, Randomization, and Combinatorial Optimization; Proceedings of the 16th APPROX Workshop; LNCS 7408* (A. Gupta, K. Jansen, J. Rolim, R. Servedio, eds.), Springer, Berlin 2012, pp. 243–253
- [51] Lovász, L. [1976]: On some connectivity properties of Eulerian graphs. *Acta Mathematica Academiae Scientiarum Hungaricae* 28 (1976), 129–138
- [52] Mitchell, J. [1999]: Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, k -MST, and related problems. *SIAM Journal on Computing* 28 (1999), 1298–1309
- [53] Mömke, T., and Svensson, O. [2011]: Approximating graphic TSP by matchings. *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science (FOCS 2011)*, 560–569
- [54] Monma, C.L., Munson, B.S., and Pulleyblank, W.R. [1990]: Minimum-weight two-connected spanning networks. *Mathematical Programming* 46 (1990), 153–171
- [55] Mucha, M. [2012]: $\frac{13}{9}$ -approximation for graphic TSP. *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science (2012)*, 30–41
- [56] Oveis Gharan, S., Saberi, A., and Singh, M. [2011]: A randomized rounding approach to the traveling salesman problem. *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2011)*, 550–559
- [57] Papadimitriou, C.H., and Vempala, S. [2006]: On the approximability of the traveling salesman problem. *Combinatorica* 26 (2006), 101–120
- [58] Papadimitriou, C.H., and Yannakakis, M. [1993]: The traveling salesman problem with distances one and two. *Mathematics of Operations Research* 18 (1993), 1–12
- [59] Qian, J., Schalekamp, F., Williamson, D.P., and van Zuylen, A. [2012]: On the integrality gap of the subtour LP for the 1,2-TSP. In: *LATIN 2012: Theoretical Informatics; LNCS 7256* (D. Fernández-Baca, ed.), Springer, Berlin 2012, pp. 606–617
- [60] Rado, R. [1942]: A theorem on independence relations. *Quarterly Journal of Mathematics* 13 (1942), 83–89
- [61] Rao, S.B., and Smith, W.D. [1998]: Approximating geometric graphs via “spanners” and “banyans”. *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC 1998)*, 540–550
- [62] Sahni, S., and Gonzalez, T. [1976]: P -complete approximation problems. *Journal of the ACM* 23 (1976), 555–565
- [63] Schalekamp, F., Williamson, D.P., and van Zuylen, A. [2012]: A proof of the Boyd-Carr conjecture. *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2012)*, pp. 1477–1486
- [64] Schrijver, A. [2003]: *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, Berlin 2003
- [65] Sebő, A. [2012]: Eight fifth approximation for TSP paths. arXiv:1209.3523; accepted for IPCO 2013
- [66] Sebő, A., and Vygen, J. [2012]: Shorter tours by nicer ears: $7/5$ -approximation for graphic TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. arXiv:1201.1870
- [67] Vishnoi, N.K. [2012]: A permanent approach to the traveling salesman problem. *Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2012)*, 76–80
- [68] Whitney, H. [1932]: Non-separable and planar graphs. *Transactions of the American Mathematical Society* 34 (1932), 339–362
- [69] Wolsey, L.A. [1980]: Heuristic analysis, linear programming and branch and bound. *Mathematical Programming Study* 13 (1980), 121–134

Discussion Column

Michel X. Goemans

Thinness Spurs Progress

In his article, Jens Vygen has discussed at length the recent developments that have occurred in the last two to three years in the world of approximation algorithms for the Traveling Salesman Problem, both in the symmetric and asymmetric case. In this discussion column, let me simply add a few comments mostly for the asymmetric case.

I will start by addressing some of the criticisms that have been directed towards approximation algorithms. One may wonder whether a real traveling salesman or the designer of circuit boards or a transportation company which routinely solves TSP problems would be satisfied or happy with tours 33.3% more costly than optimal in the symmetric case, or even a much larger factor in the asymmetric case. I doubt so, and I wouldn't either if I was in the salesman's shoes (although I have to confess that, as a traveling professor, I tremendously enjoy taking detours...). However, the renewed interest and the focus in the last few years as discussed in Vygen's article have been on understanding how to derive good tours from the optimum solution of the Held-Karp lower bound (obtained by optimizing over the symmetric or asymmetric subtour polytope), and I think this is a very important endeavor. The Held-Karp lower bound has been known to be very close to optimal (both for symmetric and asymmetric instances) and being able to round well such solutions to tours may have not just an impact for theoreticians interested in the approximability of the problem but also to practitioners. At a recent workshop devoted to the TSP in Corsica, Denis Naddef mentioned open problems he would like to be solved before he retires. Before my own retirement, I would love to see an algorithm which efficiently rounds the solution to the Held-Karp bound and produces a tour which on most typical instances is within 1 or 2% of optimal and never more than $4/3$ of the optimal (without simply running in parallel two algorithms and taking the best solution...). We are not there yet (no definite retirement plans yet...), but I am happy to see a lot of young talent getting interested in the problem. I personally have been interested in the Held-Karp lower bound since learning about its conjectured worst case of $\frac{4}{3}$ in a class taught by Laurence Wolsey in the mid 80's (while an undergraduate at UCL in Belgium) and from David Shmoys a few years later (while a graduate student at MIT). In Wolsey's paper (77) on using linear programming for the analysis of heuristics in which he proves that the integrality gap for the Held-Karp lower bound is at most $3/2$, he raises the issue of the exact worst-case gap and mentions that, at that time (in 1980), the worst gap he is aware of is $8/7$.

As emerges from Vygen's article, the worst-case quality of the Held-Karp lower bound is much more elusive in the asymmetric case (ATSP) than in the symmetric case, and so is the approximability of the associated problems. Charikar, Karloff and I (73) have constructed a family of ATSP instances for which the integrality gap is arbitrarily close to 2, but interestingly, these instances have a number of vertices exponential in $\frac{1}{\epsilon}$ to achieve a gap of $2 - \epsilon$. In contrast, in the symmetric setting, one only needs a number of vertices linear in $\frac{1}{\epsilon}$ to achieve $\frac{4}{3} - \epsilon$ for the 3-path configuration. For many other combinatorial optimization problems, this linear growth in $1/\epsilon$ is quite typical, and so this exponential growth is rather peculiar.

There is an ocean between this lower bound of 2 on the integrality for ATSP and the best known upper bound of $O(\log n / \log \log n)$ where $n = |V|$, proved recently in joint work with Asadpour et al. (71), and improving upon a long sequence of logarithmic bounds.

From this work (see Vygen's article), a key question surfaces. Is there an $\alpha > 1$ (independent of the size of the graph) for which the following holds (with the same notation as in Vygen's article)?

Given an undirected graph G and a point x in its spanning tree polytope, we can find a spanning tree T of G such that $|\delta_T(U)| \leq \alpha x(\delta(U))$ for all $U \subset V$.

This factor α is called the *thinness* of the spanning tree (with respect to x). If we were focusing on singleton cuts, we would want to construct a spanning tree satisfying some given upper bounds on its degrees, and this was the motivation behind my work on bounded-degree spanning trees (75). If we take a complete graph G on n vertices and a uniform point x (with $x_e = \frac{2}{n}$ for all edges e) then looking at singleton cuts are enough: A spanning tree with maximum degree Δ is indeed (better than) Δ -thin. Similarly if x was uniform on an expander graph, but in general, one has to consider all cuts.

A constructive/algorithmic answer to the above question (with α independent of n) would give an approximation algorithm with a constant approximation factor for the ATSP, answering a longstanding open question. How does thinness come into play in the approximability of ATSP? Although details were given in Vygen's article, let me highlight informally the key reason. One easy way to obtain an Eulerian (i.e., with equal indegree and outdegree at every vertex) directed graph is to start from a weakly connected subgraph (a spanning tree, for example) and then augment it at minimum cost into an Eulerian directed graph. Indeed, this augmentation subproblem is an easy minimum cost flow problem (see (71; 76)). However, the cost of this min-cost flow augmentation can be upper bounded linearly in terms of the thinness of the spanning tree we started from, hence the incentive for the tree to be on a diet.

Let us go back to the question of finding a thin tree. As any element x in the spanning tree polytope can be expressed (in many ways) as a convex combination of spanning trees, we could try studying the thinness of spanning trees in such a convex combination. Any convex combination can also be interpreted as a probability distribution with given marginals x . If this probability distribution is *negatively correlated* – i.e., for all $F \subset E$, we have $\mathbb{P}[F \subseteq T] \leq \prod_{e \in F} \mathbb{P}[e \in T]$ where T is a random spanning tree – then one can show (see Vygen or (71)) that the spanning tree T is $O(\log(n)/\log \log(n))$ -thin with high probability. There are several ways of obtaining a negatively correlated distribution with marginals x , and hence several ways of deriving a $O(\log(n)/\log \log(n))$ -approximation algorithm for ATSP. Asadpour et al. (71) uses a distribution with the probability of spanning tree T being proportional to $\prod_{e \in T} \lambda_e$ for a vector λ approximately equal to the optimum Lagrange multipliers for the convex program of maximizing the entropy of the distribution. Proving that this vector λ can be found efficiently is somewhat tricky. There are two other ways I am aware of that give a negatively correlated distribution with given marginals; both of these can be applied as well to any matroid polytope while the maximum entropy approach fails to give a negatively correlated distribution for some matroids. One way is mentioned in Vygen's article and uses the randomized swap rounding approach of Chekuri, Vondrák and Zenklusen (74). Yet another way is inspired by the randomized pipage rounding approach (72) as was explained to me by Rico Zenklusen. Let me describe it as it is particularly simple, has never appeared in print, and is possibly even more powerful than it appears. If one would like to express a point x in a polytope P as a convex combination of extreme points (a constructive version of Caratheodory's theorem), one possibility is to choose an arbitrary nonzero direction d within the minimal face F_x containing x and shoot in the directions d and $-d$ until one would leave P , thereby obtaining two points y and z whose corresponding minimal faces F_y and F_z have smaller dimensions than F_x . The point x can then be expressed as a convex combination of y

and z and we can recurse (here, we do not care that such a convex combination may have exponentially many terms). This is classical. But instead of choosing an arbitrary direction within F_x , we can always impose that d is one of the edges of F_x , and thus of P . In the case of the spanning tree polytope (or any matroid base polytope), the edges of the polytope have all their components 0, except for one $+1$ and one -1 . If we always use such a direction, one can prove by induction that the resulting probability distribution (convex combination) is negatively correlated: If one has negatively correlated distributions for y and z and they differ in only two components then the resulting distribution for x is negatively correlated.

The notion of thinness discussed above allows the derivation of an $O(\log(n)/\log \log(n))$ bound but may not be the easiest to handle for further, significant improvements. Indeed, computing the thinness of a tree amounts to solving a sparsest cut problem, for which the current best approximation algorithm has an approximation factor of $O(\sqrt{\log n})$ (2012 Fulkerson prize winning paper (70)). One might hope to generate a thin tree along with a certificate of thinness and get around this (current) approximability limitation, but this might be too much to ask. But there are other – this time, well-characterized – notions of thinness that may lead to constant guarantees for ATSP. For brevity (and to avoid further delaying this issue of *Optima* ...), let me mention only one of them, that comes directly from the work in (71). Instead of considering an undirected solution x (in the spanning tree polytope or a feasible solution to the symmetric subtour polytope), consider a feasible solution γ to the asymmetric subtour polytope with support A . A weakly connected set $T \subseteq A$ of arcs (or, if minimal, a weakly connected tree) is directed α -thin with respect to γ if $|\delta_T^+(U)| - |\delta_T^-(U)| \leq \alpha[\gamma(\delta^+(U)) + \gamma(\delta^-(U))]$ for all $U \subset V$. On the left-hand-side, we have the imbalance of a cut in our tree while on the right-hand-side we have the total value of γ in the cut defined by U . This directed thinness is a weaker notion (so it should be easier to get thinner trees) and has the main advantage that this directed thinness can be computed in polynomial time as a network flow problem. Furthermore, being Eulerian, the solution γ really behaves like a symmetric solution: $\gamma(\delta^+(U)) = \gamma(\delta^-(U))$ and all properties of undirected cuts (cactus and polygon representation of minimum and approximately minimum cuts, splitting-off, etc.) still hold. Producing a directed α -thin tree directly (without going through the undirected notion) appears to be challenging but well worth the effort as a constant α would give a constant approximation factor for ATSP.

In conclusion, I hope that this recent progress regarding the TSP will lead to unlocking the secrets of the quality of the Held-Karp lower bound for both the symmetric and asymmetric cases.

Michel X. Goemans, Department of Mathematics, Cambridge, MA 02139, USA. goemans@math.mit.edu

References

- [70] S. Arora, S. Rao and U. Vazirani, "Expander Flows, Geometric Embeddings and Graph Partitioning", *J. ACM*, **56**, 2009.
- [71] A. Asadpour, M.X. Goemans, A. Madry, S. Oveis Gharan and A. Saberi, "An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem", *21st ACM-SIAM Symposium on Discrete Algorithms*, 2010.
- [72] G. Calinescu, C. Chekuri, M. Pal and J. Vondrák, "Maximizing a Monotone Submodular Function subject to a Matroid Constraint", *SICOMP*, **40**, 1740–1766, 2011.
- [73] M. Charikar, M.X. Goemans and H. Karloff, "On the Integrality for the Asymmetric Traveling Salesman Problem", *Mathematics of Operations Research*, **31**, 245–252, 2006.
- [74] C. Chekuri, J. Vondrák and R. Zenklusen, "Dependent Randomized Rounding for Matroid Polytopes and Applications", 2009.
- [75] M.X. Goemans, "Minimum Bounded Degree Spanning Trees", *47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, 273–282, 2006.

- [76] M.X. Goemans, N.J.A. Harvey, K. Jain and M. Singh, “A Randomized Rounding Algorithm for the Asymmetric Traveling Salesman Problem”, arXiv:0909.0941 [cs.DS], 2009.
- [77] L.A. Wolsey, “Heuristic Analysis, Linear Programming and Branch and Bound”, *Mathematical Programming Study*, **13**, 121–134, 1980.

Prizes awarded at ISMP 2012

The Dantzig Prize

Committee: John Birge (Chair), Gerard Cornuejols, Yuri Nesterov, Eva Tardos

The George B. Dantzig prize is awarded for original research, which by its originality, breadth, and scope, is having a major impact on the field of mathematical programming. The 2012 winners are Professors

1. *Jorge Nocedal* from Northwestern University and
2. *Laurence Wolsey* from Université Catholique de Louvain (UCL).

Jorge Nocedal has made fundamental contributions to the theory of nonlinear optimization methods, has developed new algorithms that expand the range of efficiently solvable optimization models, and has created widely distributed software that enables new developments in numerous application areas. Laurence Wolsey has been one of the most influential scholars in the field of mathematical optimization, contributing significantly to foundational understanding of the geometry of mixed integer programs, to duality theory in discrete optimization, and to the development of effective new methods for a variety of applications, particularly in production planning and scheduling.

The Beale-Orchard-Hays Prize

Committee: M. Ferris (Chair), P. Gill, T. Kelley, J. Lee.

In a unanimous decision, the selection committee for the Beale Orchard Hays Prize for 2012 decided that the award be given to *Michael Grant and Stephen Boyd* for the software CVX as described in the following two papers:

1. CVX: Matlab software for disciplined convex programming, version 1.21, <http://cvxr.com/cvx>, April 2011.
2. “Graph Implementations for Nonsmooth Convex Programs” in *Recent Advances in Learning and Control*, V. Blondel, S. Boyd and H. Kimura (eds), pp. 95–110, *Lecture Notes in Control and Informational Sciences*, Springer, 2008.

The nomination states that “CVX is a modeling language for convex programming that has been implemented in Matlab. It also automatically links to semidefinite solvers ... CVX makes convex programming as easy as Matlab makes matrix computation.” The committee feels that this work is very well respected within our community and is used extensively for both research and teaching at a number of high-profile institutions. In particular, it provides a unique, well-documented tool for prototyping and exploring existing and emerging applications of convex optimization.

The Fulkerson Prize

Committee: K. Aardal (Chair), Paul Seymour, Richard Stanley.

The award was given to three papers:

1. Sanjeev Arora, Satish Rao, and Umesh Vazirani, “Expander flows, geometric embeddings and graph partitioning”, *J. ACM*, **56** (2009), 1–37.
2. Anders Johansson, Jeff Kahn, and Van Vu, “Factors in random graphs”, *Random Structures and Algorithms* **33** (2008), 1–28.
3. László Lovász and Balázs Szegedy, “Limits of dense graph sequences”, *Journal of Combinatorial Theory Series B* **96** (2006), 933–957.

Citations can be found at <http://www.mathopt.org/?nav=fulkerson>.

The Lagrange Prize in Continuous Optimization

Committee: T. Terlaky (Chair), K. Anstreicher, D. Goldfarb, T. Liebling.

The Lagrange Prize was awarded to *Emmanuel J. Candès and Benjamin Recht* for their paper “Exact matrix completion via convex optimization”, *Foundations of Computational Mathematics* **9** (2009), 717–772.

The paper of Candès and Recht was selected because of its exposition excellence, the current importance of the topic and the impressive number of citations in three years. It also opens Semidefinite Optimization to a fascinating new field of applications and introduces a very clever mathematical approach for proving probabilistic tractability of certain NP hard problems.

Paul Y. Tseng Memorial Lectureship in Continuous Optimization

Committee: S. Leyffer (Chair), D. Li, S. Ulbrich, N. Xiu.

The first recipient of Paul Y. Tseng memorial lectureship was *Yinyu Ye* of Stanford University. Yinyu Ye has been at the forefront of continuous optimization and in particular research into interior-point methods for over 20 years. His accomplishments span the breadth of optimization including fundamental theoretical contributions into interior-point methods, the development of semi-definite programming software, and the promotion of novel applications of optimizations in economic markets and distance geometry.

In addition, Yinyu Ye has played a pivotal role in promoting optimization research in the Asia-Pacific region. He was a founding editor of the *Pacific Journal of Optimization* and he has held numerous honorary appointments at Chinese Universities, where he regularly teaches popular tutorial lectures and supervises students.

A. W. Tucker Prize

Committee: D. Ralph (Chair), M. Anjos, F. Eisenbrand, B. Fortz, B. Morini.

The Tucker Prize for an outstanding doctoral thesis has been awarded to *Oliver Friedmann*, Department of Computer Science, Ludwig-Maximilians-Universität in Munich, Germany, for his thesis “Exponential Lower Bounds for Solving Infinitary Payoff Games and Linear Programs.”

One of the most prominent mysteries in Optimization is the question of whether a linear program can be solved in strongly-polynomial time. A strongly polynomial-time method would be polynomial in the dimension and in the number of inequalities only, whereas the complexity of the known weakly-polynomial time algorithms for linear programming, like the ellipsoid method or variants of the interior-point method, also depend on the binary encoding length of the input. The simplex method, though one of the oldest methods for linear programming, still is a candidate for such a strongly polynomial time algorithm. This would require the existence of a pivoting rule that results in a polynomial number of pivot steps. Since the famous Klee-Minty example, many techniques for deriving exponential lower bounds on the number of iterations for particular pivoting rules have been found.

Some very important pivoting rules, however, have resisted a super-polynomial lower-bound proof for a very long time. Among them the Random-Facet pivoting rule and Zadeh’s pivoting rule. Random-Facet has been shown to yield sub-exponential running time of the simplex method independently by Kalai as well as by Matousek, Sharir and Welzl.

Zadeh was a postdoc at Stanford in 1980, when he published a technical report with his least-entered rule: enter the improving variable that has been entered least often. In a hand-written letter

to Viktor Klee he offered \$1000 to the person who either showed this rule to be a polynomial pivoting rule for the simplex method, or provided a counterexample to it being a polynomial method. Consequently, Zadeh's rule is very well known in the linear-programming community.

In his thesis, Oliver Friedmann has shown super-polynomial lower bounds for pivoting rules in a groundbreaking way. The novelty of his approach is to establish a connection from policy iteration for 2-player parity games and Markov decision processes to pivoting in linear programs. In his paper *Subexponential lower bounds for randomized pivoting rules for solving linear programs*, coauthored with Hansen and Zwick (Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC'11, San Jose, CA, USA, 2011), Friedmann shows a super-polynomial bound on the Random-Facet pivoting rule. This paper was awarded the prestigious STOC best paper award. This line of work, initiated by Friedmann, shows that the standard strategy iteration algorithm for parity games may require an exponential number of iterations. By giving analogous results for Markov decision processes, Friedmann extends super-polynomial lower bounds to pivoting in linear programming.

The thesis of Friedmann lays out this connection of improvement strategies for games and pivoting. Two of the most prominent results are the aforementioned lower bounds for Random-Facet and Zadeh's rule. But, with this new connection, other pivoting rules that resisted super-polynomial lower-bound proofs have also been shown to be non-polynomial, like the Random-Edge rule and, in a recent publication of Friedmann, Cunningham's rule as well.

Oliver Friedmann is 27 years old (at the date of receiving the Tucker Prize). His undergraduate, master's-level and Ph.D. degrees were all undertaken in the Department of Computer Science at the Ludwig-Maximilians-Universität in Munich, and were completed in 2006, 2008 and 2011 respectively. His Ph.D. thesis was completed in only 2.5 years under supervision from Martin Hofmann and Martin Lange.

With his thesis, Friedmann has built bridges between so-far seemingly unrelated fields, enriched optimization with novel ideas and techniques and achieved groundbreaking results that settled

many longstanding open problems. This thesis truly deserves to be awarded the Tucker Prize 2012.

The other two Tucker Prize finalists chosen by this year's Tucker Prize Committee are *Amitabh Basu* and *Guanghai Lan*.

Amitabh Basu obtained his undergraduate degree in Computer Science and Engineering from the Indian Institute of Technology, Delhi in 2004, and received an M.S. in Computer Science from Stony Brook University in 2006. In May 2010, he finished a Ph.D. in Algorithms, Combinatorics and Optimization from the Tepper School of Business, Carnegie Mellon University advised by Gerard Cornuéjols. He is currently a visiting assistant professor in the Department of Mathematics at the University of California, Davis. The thesis of Basu is entitled "Corner Polyhedra and Maximal Lattice-free Convex Sets: A Geometric Approach to Cutting Planes".

Guanghai Lan obtained his undergraduate degree in Mechanical Engineering from the Xiangtan University, China, in 1996 and went on to complete two master's degrees, one in Mechanical Engineering at the Shanghai Jiao Tong University, China, 1999, and the other in Industrial Engineering at the University of Louisville, Kentucky, 2004. In January 2009 he completed his Ph.D. in Industrial and Systems Engineering at Georgia Institute of Technology supervised by Arkadi Nemirovski and co-advised by Renato Monteiro and Alexander Shapiro. He is currently an Assistant Professor of Industrial and Systems Engineering at the University of Florida. Lan's dissertation is entitled "Convex Optimization under Inexact First-Order Information, concerns the design and complexity analysis of first-order methods for solving convex optimization problems under a stochastic oracle".

The Tucker Prize Committee was both humbled and inspired the large number of outstanding doctoral theses submitted. Beyond the three finalists we would like to unofficially commend three nominees for their truly superb work (in alphabetical order): 1. João Gouveia for his 2011 PhD thesis "Geometry of Sums of Squares Relaxations" at the University of Washington; 2. Fernando de Oliveira Filho for his 2009 PhD thesis "New Bounds for Geometric Packing and Coloring via Harmonic Analysis and Optimization" at the University of Amsterdam; and Neil Olver for his 2010 PhD thesis "Robust Network Design" at McGill University.



Many members of MOS may be aware that our Society's Chair, Philippe Toint, turned 60 in April. As part of a year of celebrations, September's IMA Conference on Numerical Linear Algebra and Optimisation in Birmingham, UK, included a session dedicated to Philippe's achievements. This was followed by a mass assembly of Philippe Toint impersonators ... see if you can spot the original! Orchestration by Cordia Cartis and Nick Gould, photograph by Michal Kocvara. The celebrations will continue in Toulouse next July, see <http://www.fondation-stae.net/fr/optimization-july2013.html>.

Springer Book Archives

Make your book part of this archiving initiative!

Have you already heard about the Springer Book Archives? On this page you will find all relevant information to make sure that your book will be included in this major initiative.

At this point you may be wondering what the Springer Book Archives are all about. The archive will consist of high-quality digital versions of books from Springer (and about 40 imprints) published between 1842 and 2005. It will include many titles that are now out-of-print and make them available again, both online and in print.

This project affirms our commitment to preserving valuable scholarly content. It will enable researchers all over the world to access a wealth of information through their libraries. Your book plays an important part in this effort!

How to include my work in the Springer Book Archives?

Please go to the URL

<http://www.springer.com/authors/oba?SGWID=0-1726313-0-0-0>
and follow the hints given there.

IPCO 2013

The 16th Conference on Integer Programming and Combinatorial Optimization (IPCO XVI) will be held March 18–20, 2013, at Universidad Tecnica Federico Santa Maria in Valparaiso, Chile.

The IPCO conference is held every year, except for those in which the International Symposium on Mathematical Programming takes place. The conference is a forum for researchers and practitioners

working on various aspects of integer programming and combinatorial optimization. The aim is to present recent developments in theory, computation, and applications. The scope of IPCO is viewed in a broad sense, to include algorithmic and structural results in integer programming and combinatorial optimization as well as revealing computational studies and novel applications of discrete optimization to practical problems.

Program committee: Chandra Chekuri (UIUC), Bill Cook (Georgia Tech), Jose Correa (U Chile, local arrangements chair), Jesus De Loera (UC Davis), Michel Goemans (MIT, chair), Volker Kaibel (U Magdeburg), Jon Lee (U Michigan), Francois Margot (CMU), Thomas McCormick (UBC), Andreas Schulz (MIT), David Shmoys (Cornell), Zoltan Szigeti (INP Grenoble), Robert Weismantel (ETH), Giacomo Zambelli (LSE)

Information: <http://ipco2013.dim.uchile.cl>

MIP 2013

The MIP2013 workshop on Mixed Integer Programming will take place in Madison, Wisconsin, at the Wisconsin Institutes for Discovery from July 22 to July 25, 2013. MIP2013 will be the tenth in a series of annual workshops held in North America designed to bring the integer programming community together to discuss recent developments in the field.

The program committee for the workshop consists of Amitabh Basu (UC Davis), Dan Bienstock (Columbia), Alberto Del Pia (ETH Zurich), Santanu Dey (Georgia Tech), and Jim Luedtke (UW Madison). Michael Ferris, Jeff Linderoth, Jim Luedtke, and Christos Maravelias make up the local organizing committee.

As more information is available about the workshop, it will be made available at <https://events.discovery.wisc.edu/mip2013/>.

Application for Membership

I wish to enroll as a member of the Society. My subscription is for my personal use and not for the benefit of any library or institution.

- I will pay my membership dues on receipt of your invoice.
 I wish to pay by credit card (Master/Euro or Visa).

Credit card no.

Expiration date

Family name

Mailing address

Telephone no.

Telefax no.

E-mail

Signature

Mail to:

Mathematical Optimization Society
3600 Market St, 6th Floor
Philadelphia, PA 19104-2688
USA

Cheques or money orders should be made payable to The Mathematical Optimization Society, Inc. Dues for 2012, including subscription to the journal *Mathematical Programming*, are US \$90. Retired are \$45. Student applications: Dues are \$22.50. Have a faculty member verify your student status and send application with dues to above address.

Faculty verifying status

Institution

IMPRINT

Editor: Katya Scheinberg, Department of Industrial and Systems Engineering, Lehigh University, Bethlehem, PA, USA. katyas@lehigh.edu ■ **Co-Editors:** Samuel Burer, Department of Management Sciences, The University of Iowa, Iowa City, IA 52242-1994, USA. samuel-burer@uiowa.edu ■ Volker Kaibel, Institut für Mathematische Optimierung, Otto-von-Guericke Universität Magdeburg, Universitätsplatz 2, 39108 Magdeburg, Germany. kaibel@ovgu.de ■ **Founding Editor:** Donald W. Hearn ■ Published by the Mathematical Optimization Society ■ Design and typesetting by Christoph Eyrich, Berlin, Germany. optima@0x45.de ■ Printed by Oktoberdruck AG, Berlin, Germany.