Operations Research Letters 40 (2012) 36–38

Contents lists available at SciVerse ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Multi-criteria TSP: Min and max combined

Bodo Manthey

University of Twente, Department of Applied Mathematics, P.O. Box 217, 7500 AE Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 13 July 2011
Accepted 11 October 2011
Available online 26 October 2011

Keywords:
Multi-criteria optimization
Pareto optimization
Approximation algorithms
Traveling salesman problem

ABSTRACT

We present randomized approximation algorithms for multi-criteria traveling salesman problems (TSP), where some objective functions should be minimized while others should be maximized. For the symmetric multi-criteria TSP (STSP), we present an algorithm that computes $(2/3, 3+\varepsilon)$-approximate Pareto curves. Here, the first parameter is the approximation ratio for the objectives that should be maximized, and the second parameter is the ratio for the objectives that should be minimized. For the asymmetric multi-criteria TSP (ATSP), we obtain an approximation performance of $(1/2, \log_2 n+\varepsilon)$.

1. Multi-criteria TSP

The traveling salesman problem (TSP) is a basic optimization problem. The goal is to find Hamiltonian cycles of maximum or minimum weight. An instance of *Max-TSP* is a complete graph $G=(V,E)$ with edge weights $w: E \to \mathbb{Q}_+$. The goal is to find a *Hamiltonian cycle* (also called a *tour*) of maximum weight, where the weight of a Hamiltonian cycle is the sum of its edge weights. (The weight of an arbitrary set of edges is analogously defined.) If G is undirected, then we speak of *Max-STSP* (symmetric TSP). If G is directed, we have *Max-ATSP* (asymmetric TSP). *Min-TSP* is similarly defined, but now the edge weights $d: E \to \mathbb{Q}_+$ are required to fulfill the triangle inequality: $d(u,v) \leq d(u,x)+d(x,v)$ for all $u, v, x \in V$. (Without the triangle inequality, Min-TSP does not allow for any approximation [2], and we are concerned with approximation algorithms.) The aim is to find a Hamiltonian cycle of minimum weight. *Min-STSP* is the symmetric variant, where G is undirected, while *Min-ATSP* is the asymmetric variant.

However, we often have more than one objective function: we might want to minimize travel time, travel expenses, etc., while maximizing, e.g., our profit along the way. This gives rise to multi-criteria TSP. So far, multi-criteria TSP has only been considered in the setting where either all objectives should be minimized or all objectives should be maximized. In this paper, we consider the general setting with both types of objectives simultaneously.

If k objectives are to be maximized and ℓ objectives are to be minimized, then we have *k-Max-ℓ-Min-ATSP* and *k-Max-ℓ-Min-STSP*. For $\ell=0$ or $k=0$, we have *k-Max-ATSP* and *k-Max-STSP* as well as *ℓ-Min-ATSP* and *ℓ-Min-STSP*, respectively.

Unfortunately, if we have more than one objective function, there is no natural notion of an optimal solution. Instead, we have to content ourselves with trade-off solutions. Therefore, the notion of *Pareto curves* has been introduced (cf. [5]): a Pareto curve is a set of solutions that are optimal trade-offs between the different objective functions.

Let us describe this more formally. An instance of k-Max-ℓ-Min-ATSP (or k-Max-ℓ-Min-STSP) is a directed (undirected) complete graph $G=(V,E)$ with edge weights $w_1, \ldots, w_k : E \to \mathbb{Q}_+$ and $d_1, \ldots, d_\ell : E \to \mathbb{Q}_+$, where each d_i satisfies the triangle inequality. The functions $w_1, \ldots, w_k$ should be maximized while $d_1, \ldots, d_\ell$ should be minimized. We call $w_1, \ldots, w_k$ the *max objectives* and $d_1, \ldots, d_\ell$ the *min objectives*. For convenience, let $w=(w_1, \ldots, w_k)$ and $d=(d_1, \ldots, d_\ell)$. Inequalities of vectors are meant component-wise. A Hamiltonian cycle H *dominates* another Hamiltonian cycle H' if $w(H) \geq w(H')$ and $d(H) \leq d(H')$ and at least one of these inequalities is strict. This means that H is strictly preferable to H'. A *Pareto curve* contains all solutions that are not dominated by another solution. For other optimization problems, this is defined analogously. Unfortunately, since Pareto curves for the TSP cannot be computed efficiently (because, first, they can be of exponential size and, second, we inherit the approximation hardness of the TSP with a single objective function), we have to be satisfied with approximate Pareto curves.

A set $\mathcal{P}$ of Hamiltonian cycles is called an $(\alpha_{\max}, \alpha_{\min})$-*approximate Pareto curve* for the instance (G, w, d) if the following holds: for every Hamiltonian cycle $\tilde{H}$ of G, there exists a Hamiltonian cycle $H \in \mathcal{P}$ of G with $w(H) \geq \alpha_{\max} w(\tilde{H})$ and $d(H) \leq \alpha_{\min} d(\tilde{H})$. We have $\alpha_{\max} \leq 1$, $\alpha_{\min} \geq 1$, and a $(1,1)$-approximate Pareto curve is a Pareto curve.

An algorithm is called an *$(\alpha_{\max}, \alpha_{\min})$-approximation algorithm* if, given G, w, and d, it computes an $(\alpha_{\max}, \alpha_{\min})$-approximate Pareto curve. It is called a *randomized $(\alpha_{\max}, \alpha_{\min})$-approximation* if its success probability is at least $1/2$. This success probability can be amplified by executing the algorithm several times.

E-mail address: b.manthey@utwente.nl.

0167-6377/$ – see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.orl.2011.10.012

Table 1
Known approximation ratios for TSP and the results of this paper.

Variant	Single-criterion	Multi-criteria	Combined
Max-STSP	7/9 [13]	2/3 [7]	$(2/3, 3+\varepsilon)$
Min-STSP	3/2 [2]	$2+\varepsilon$ [12]	
Max-ATSP	2/3 [8]	1/2 [7]	$(1/2, \log_2 n+\varepsilon)$
Min-ATSP	$O\left(\frac{\log n}{\log\log n}\right)$ [1]	$\log_2 n+\varepsilon$ [10]	

A *fully polynomial time approximation scheme* (FPTAS) for a multi-criteria optimization problem computes $(1-\varepsilon, 1+\varepsilon)$-approximate Pareto curves in time polynomial in the size of the instance and $1/\varepsilon$ for all $\varepsilon > 0$. Multi-criteria matching (computing a perfect of minimum or maximum weight or computing a maximum-weight matching that is not necessarily perfect) admits a *randomized FPTAS* [14], i. e., the algorithm succeeds in computing a $(1-\varepsilon, 1+\varepsilon)$-approximate Pareto curve with a probability of at least 1/2. A *cycle cover* of a graph is a set of vertex-disjoint cycles such that every vertex is part of exactly one cycle. The randomized FPTAS for matching yields also a randomized FPTAS for the multi-criteria cycle cover problem [12]. A *path cover* of a graph is a set of vertex-disjoint paths such that every vertex is part of exactly one path.

Known results. ℓ-Min-STSP admits a $(2+\varepsilon)$-approximation algorithm [12]. For the special case of $\ell = 2$, there even is a 2 approximation [6]. ℓ-Min-ATSP can be approximated with a factor of $\log_2 n + \varepsilon$ [9]. k-Max-ATSP and k-Max-STSP have first been studied by Bläser et al. [4], who have achieved approximation ratios of $\frac{1}{k+1} - \varepsilon$ and $\frac{1}{k} - \varepsilon$, respectively. This has subsequently been improved to 1/2 and 2/3, respectively [7]. Table 1 summarizes the known results and compares them to the approximation ratios obtained for single-criterion TSP.

As far as we are aware, nothing is known yet about the approximability of multi-criteria TSP with both min and max objectives. In fact, with only a few exceptions [3,14], little is known about the approximability of any multi-criteria optimization problem with both min and max objectives.

New results. We present a generic algorithm for computing approximate Pareto curves for k-Max-ℓ-Min-ATSP and k-Max-ℓ-Min-STSP. Our algorithm achieves approximation ratios of $(1/2, \log_2 n + \varepsilon)$ and $(2/3, 3+\varepsilon)$, respectively. (See also Table 1.) The running-time of our algorithms is polynomial for any fixed ℓ, k, and ε. Note that an exponential dependence on k and ℓ is unavoidable, as the size of even approximate Pareto curves can increase exponentially with the number of objective functions.

The main idea of our algorithms is to first find a collection of paths that have sufficient weight for the max objectives. This is basically done using a (slightly modified) approximation algorithm for multi-criteria Max-TSP, while we have to take into account that these paths are not too heavy with respect to d. After that, we connect these paths to get Hamiltonian cycles using an approximation algorithm for multi-criteria Min-TSP. In this second step, we only pay attention to the min objectives; we already have enough weight for the max objectives, and adding further edges does not decrease the weight. (A preliminary version of this paper has been presented at the 7th Int. Workshop on Approximation and Online Algorithms [11]. The present paper improves the results presented there and simplifies the proofs.)

2. Approximation algorithm for min and max

Now we present and analyze our framework for approximation algorithms for multi-criteria TSP that will yield the results mentioned above. The framework will be for both STSP and ATSP.

For our framework to work, we need the concept of a canonical algorithm. An algorithm is called an (α, ε)-*canonical algorithm* if

$\mathcal{P}_{\text{TSP}} \leftarrow$ TSP-Approx(G, w, d, ε)
input: (un)directed complete graph $G = (V, E)$, $w : E \to \mathbb{Q}_+^k$, $d : E \to \mathbb{Q}_+^\ell$, $\varepsilon > 0$
output: $(\alpha, r_{n/2} + 1 + \varepsilon)$-approximate Pareto curve with probability at least 1/2
$\mathcal{P}_{\text{paths}} \leftarrow$ GRW-Approx$(G, w, d, \varepsilon/2)$
for all $P \in \mathcal{P}_{\text{paths}}$ **do**
  let V' be the start-points of paths of P
  $\mathcal{P}_{\text{TSP}}{}' \leftarrow$ MinTSP-Approx$(V', d, \varepsilon/2)$
  for all $H' \in \mathcal{P}_{\text{TSP}}{}'$ **do**
    construct a tour H from $P \cup H'$ as in the proof of Theorem 1
    add H to $\mathcal{P}_{\text{TSP}}$

Algorithm 1: Approximation algorithm for k-Max-ℓ-Min-ATSP and k-Max-ℓ-Min-STSP.

it has the following property: on input $(G = (V, E), d, w)$, it computes a set of path covers such that for each Hamiltonian cycle $\tilde{H}$ of G it outputs a path cover P of G with the following properties. First, by adding an edge from the end-point of each path in P to its corresponding start-point, one obtains a cycle cover $C \supseteq P$ such that $d(C) \leq (1+\varepsilon) \cdot d(\tilde{H})$. Second, $w(P) \geq \alpha w(\tilde{H})$. Third, (V, P) consist of at most $n/2$ paths.

Let us now describe how the algorithm of Glaßer et al. [7] for k-Max-STSP and k-Max-ATSP can be turned into an $(2/3, \varepsilon)$ and $(1/2, \varepsilon)$-canonical algorithm, respectively, for any $\varepsilon > 0$ for the symmetric and asymmetric variant. First, instead of computing $\left(1 - \frac{1}{n}\right)$-approximate Pareto curves of cycle covers with respect to w (Glaßer et al. compute *edge-fixed cycle covers*, but this does not make a difference: we can fix some edges by removing them and the corresponding vertices from the graph), we compute $\left(1 - \frac{1}{n}, 1+\varepsilon\right)$-approximate Pareto curves of cycle covers with respect to both w and d (using Papadimitriou and Yannakakis' randomized FPTAS [14]). We proceed with the original algorithm, ignoring d. From every such cycle cover C, the algorithm constructs a path cover P by removing exactly one edge from each cycle. If C is a cycle cover with $w(C) \geq \left(1 - \frac{1}{n}\right) \cdot w(\tilde{H})$ for some tour $\tilde{H}$ and C contains the right collection of fixed edges, then $w(P) \geq \frac{2}{3} \cdot w(\tilde{H})$ (in the symmetric case) or $w(P) \geq \frac{1}{2} \cdot w(\tilde{H})$ (in the asymmetric case). (In order to find such a cycle cover C, we have to fix the heavy edges of the cycle cover in advance. Since there are always only a few heavy edges, their algorithm iterates over all possibilities while maintaining polynomial running-time. Thus, we are guaranteed to find the right collection.) Furthermore, as we compute $\left(1 - \frac{1}{n}, 1+\varepsilon\right)$-approximate Pareto curves of cycle covers, we can achieve $d(C) \leq (1+\varepsilon) \cdot d(\tilde{H})$. The graph (V, P) consists of at most $n/2$ paths. Finally, we remove the last line of the algorithm [7], where the paths are connected to a Hamiltonian cycle in an arbitrary way and simply output the path cover. We denote the canonical algorithm obtained from the algorithm by Glaßer et al. by GRW-Approx.

Theorem 1. *Assume that there is an (α, ε)-canonical randomized polynomial-time algorithm for k-Max-ℓ-Min-ATSP or k-Max-ℓ-Min-STSP for any $\varepsilon > 0$. Assume further that ℓ-Min-ATSP or ℓ-Min-STSP can be approximated with an approximation ratio of r_n (for n vertices) by a randomized polynomial-time algorithm.*

Then TSP-Approx *(Algorithm 1) is a randomized polynomial-time $(\alpha, r_{n/2} + 1 + \varepsilon)$-approximation algorithm for k-Max-ℓ-Min-ATSP or k-Max-ℓ-Min-STSP.*

Proof. Let us analyze the approximation ratio first. To do this, we assume that all randomized computations are successful. After that, we analyze success probability and running-time. Let $\tilde{H}$ be an arbitrary Hamiltonian cycle. Let GRW-Approx be (α, ε)-canonical.

We have to show that there exists a Hamiltonian cycle $H \in \mathcal{P}_{\text{TSP}}$ with $w(H) \geq \alpha \cdot w(\tilde{H})$ and $d(H) \leq (r_{n/2} + 1 + \varepsilon) \cdot d(\tilde{H})$.

The set $\mathcal{P}_{\text{paths}}$ contains a path cover $P \subseteq E$ with $d(P) \leq (1+\varepsilon) \cdot d(\tilde{H})$ and $w(P) \geq \alpha w(\tilde{H})$. Even stronger, by assumption, there is a cycle cover $C \supseteq P$ with $d(C) \leq (1+\varepsilon) \cdot d(\tilde{H})$.

By construction, all paths of the path cover P consist of at least one edge, i.e., (V, P) does not contain isolated vertices. Thus, we have $|V'| \leq |V|/2$. This implies that $\mathcal{P}_{\text{TSP}}'$ (in the corresponding iteration) contains a tour H' with $d(H') \leq r_{n/2} d(\tilde{H})$. We obtain the Hamiltonian cycle H from P and H' as follows: assume that P contains a path from u to v and H' contains an edge from u to x. Then we add the path from u to v plus the edge (v, x) to H. We do this for all paths of P. The triangle inequality guarantees $d(v, x) \leq d(v, u) + d(u, x)$. Note that the edge (v, u) is part of the corresponding cycle (that contains u and v) that is missing in P but is contained in C. This together yields

$$d(H) \leq d(H') + d(C) \leq (r_{n/2} + 1 + \varepsilon) \cdot d(\tilde{H}).$$

Furthermore, by construction, we have $H \supseteq P$. Thus, we have $w(H) \geq w(P) \geq \alpha w(\tilde{H})$.

The error probabilities of the randomized computations in lines 1 and 4 can be chosen small enough such that they sum up to at most 1/2. This yields that the probability that one of the computations fails is at most 1/2. The running-time follows because the running-times of MinATSP-Approx and GRW-Approx are polynomial if k and ℓ are fixed. □

Immediate consequences of the theorem above are the following two corollaries.

Corollary 2. *For every fixed $k, \ell \in \mathbb{N}$ and every $\varepsilon > 0$, there is a $(2/3, 3+\varepsilon)$-approximation algorithm for k-Max-ℓ-Min-STSP. Its running-time is polynomial in the input size and $1/\varepsilon$.*

Proof. By using $r_n = 2+\varepsilon$ and $\alpha = 2/3$, we obtain a $(2/3, 3+2\varepsilon)$-approximation from Theorem 1 for every $\varepsilon > 0$, which proves the assertion. □

Corollary 3. *For every fixed $k, \ell \in \mathbb{N}$ and every $\varepsilon > 0$, there is a $(1/2, \log_2 n + \varepsilon)$-approximation algorithm for k-Max-ℓ-Min-ATSP. Its running-time is polynomial in the input size and $1/\varepsilon$.*

Proof. We use $r_n = \log_2 n + \varepsilon$ and $\alpha = 1/2$. Then we obtain an approximation ratio of $(1/2, \log_2(n/2)+\varepsilon+1+\varepsilon) = (1/2, \log_2 n + 2\varepsilon)$ for every $\varepsilon > 0$. □

References

[1] Arash Asadpour, Michel X. Goemans, Aleksander Madry, Shayan Oveis Gharan, Amin Saberi, An $O(\log n/\log\log n)$-approximation algorithm for the asymmetric traveling salesman problem, in: Proc. of the 21st Ann. ACM-SIAM Symp. on Discrete Algorithms, SODA, SIAM, 2010, pp. 379–389.

[2] Giorgio Ausiello, Pierluigi Crescenzi, Giorgio Gambosi, Viggo Kann, Alberto Marchetti, Marco Protasi, Complexity and Approximation: Combinatorial Optimization Problems and Their Approximability Properties, Springer, 1999.

[3] André Berger, Vincenzo Bonifaci, Fabrizio Grandoni, Guido Schäfer, Budgeted matching and budgeted matroid intersection via the gasoline puzzle, Mathematical Programming 128 (1–2) (2011) 355–372.

[4] Markus Bläser, Bodo Manthey, Oliver Putz, Approximating multi-criteria Max-TSP, in: Proc. of the 16th Ann. European Symp. on Algorithms, ESA, in: Lecture Notes in Computer Science, vol. 5193, Springer, 2008, pp. 185–197.

[5] Matthias Ehrgott, Multicriteria Optimization, Springer, 2005.

[6] Christian Glaßer, Christian Reitwießner, Maximilian Witek, Improved and generalized approximations for two-objective traveling salesman, Report 09-076, Revision 1, Electronic Colloquium on Computational Complexity (ECCC), 2010.

[7] Christian Glaßer, Christian Reitwießner, Maximilian Witek, Applications of discrepancy theory in multiobjective approximation, in: Proc. of the 31st Conf. on Foundations of Software Technology and Theoretical Computer Science, FSTTCS, 2011 (in press).

[8] Haim Kaplan, Moshe Lewenstein, Nira Shafrir, Maxim I. Sviridenko, Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs, Journal of the ACM 52 (4) (2005) 602–626.

[9] Bodo Manthey, On approximating multi-criteria TSP, ACM Transactions on Algorithms (in press).

[10] Bodo Manthey, On approximating multi-criteria TSP, in: Proc. of the 26th Int. Symp. on Theoretical Aspects of Computer Science, STACS, 2009, pp. 637–648.

[11] Bodo Manthey, Multi-criteria TSP: min and max combined, in: Proc. of the 7th Workshop on Approximation and Online Algorithms, WAOA 2009, in: Lecture Notes in Computer Science, vol. 5893, Springer, 2010, pp. 205–216.

[12] Bodo Manthey, L. Shankar Ram, Approximation algorithms for multi-criteria traveling salesman problems, Algorithmica 53 (1) (2009) 69–88.

[13] Katarzyna Paluch, Marcin Mucha, Aleksander Madry, A 7/9 approximation algorithm for the maximum traveling salesman problem, in: Proc. of the 12th Int. Workshop on Approximation Algorithms for Combinatorial Optimization Problems, APPROX, in: Lecture Notes in Computer Science, vol. 5687, Springer, 2009, pp. 298–311.

[14] Christos H. Papadimitriou, Mihalis Yannakakis, On the approximability of trade-offs and optimal access of web sources, in: Proc. of the 41st Ann. IEEE Symp. on Foundations of Computer Science, FOCS, IEEE Computer Society, 2000, pp. 86–92.