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
b.manthey@utwente.nl

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 $(2/3-\varepsilon, 4+\varepsilon)$ 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 present an algorithm that computes $(1/2 - \varepsilon, \log_2 n + \varepsilon)$ approximate Pareto curves. In order to obtain these results, we simplify the existing approximation algorithms for multi-criteria Max-STSP and Max-ATSP. Finally, we give algorithms with improved ratios for some special cases.

1 Multi-Criteria TSP

1.1 Traveling Salesman Problems

The traveling salesman problem (TSP) is a basic problem in combinatorial optimization. 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 fulfil the triangle inequality: $d(u, v) \leq d(u, x) + d(x, v)$ for all $u, v, x \in V$ (without the triangle inequality, approximating the problem is impossible). 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.

All four variants of TSP are NP-hard and APX-hard. Thus, we are in need of approximation algorithms. Christofides' algorithm [14] achieves a ratio of 3/2 for Min-STSP. Min-ATSP can be approximated with a factor of $\frac{2}{3} \cdot \log_2 n$, where n is the number of vertices of the instance [8]. The currently best approximation algorithm for Max-STSP achieves an approximation ratio of 7/9 [12], and the currently best algorithm for Max-ATSP achieves a ratio of 2/3 [9].

Cycle covers are one of the main tools for designing approximation algorithms for the TSP [3, 8, 9, 12]. A *cycle cover* of a graph is a set of vertex-disjoint cycles such that every vertex is part of exactly one cycle. Hamiltonian cycles are special cases of cycle covers that consist just of a single cycle. The general idea is to compute an initial cycle cover, and then we join the cycles to obtain a Hamiltonian cycle.

1.2 Multi-Criteria Optimization

In many optimization problems, there is more than one objective function. This is also the case for the TSP: We might want to minimize travel time, expenses, number of flight changes, etc., while a taxi driver might want to maximize his profit, or we want to maximize, for instance, our profit along the way. This gives rise to multi-criteria TSP, where Hamiltonian cycles are sought that optimize several objectives simultaneously. However, as far as we are aware, multi-criteria TSP has only been considered in a restricted 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 at the same time.

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*. If the number of criteria does not matter, we will also speak of *MC-ATSP* and *MC-STSP*. If $\ell = 0$ or $k = 0$, then we obtain the special cases *k-Max-ATSP* and *k-Max-STSP* as well as *ℓ-Min-ATSP* and *ℓ-Min-STSP*. Analogously, if the number of criteria is unimportant, we have *MC-Max-ATSP* and so on.

With respect to a single objective function, the notion of an optimal solution is well-defined. But if we have more than one objective function, there is no natural notion of a best choice. Instead, we have to content ourselves with trade-off solutions. The goal of multi-criteria optimization is to deal with this dilemma. In order to transfer the notion of an optimal solutions to multi-criteria optimization problems, *Pareto curves* (also known as *Pareto sets* or *efficient sets*) were introduced introduced (cf. Ehrgott [6]). A Pareto curve is a set of solutions that are potential optimal choices.

An instance of k-Max-ℓ-Min-ATSP is a directed complete graph $G = (V, E)$ with edge weights $w_1, \ldots, w_k : E \to \mathbb{Q}_+$ and $d_1, \ldots, d_\ell : E \to \mathbb{Q}_+$. 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* of solutions contains all solutions that are not dominated by another solution. For other optimization problems, multi-criteria variants are defined analogously.

Unfortunately, Pareto curves cannot be computed efficiently in many cases: First, they are often of exponential size. Second, they are NP-hard to compute even for otherwise easy optimization problems. Third, TSP is NP-hard already with a single objective function, and optimization problems do not become easier with more objectives involved. Therefore, we have to be satisfied with approximate Pareto curves.

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

An algorithm is called an (α, β) *approximation algorithm* if, given G and w, it computes an (α, β) approximate Pareto curve. It is called a *randomized* (α, β) *approximation* if its success probability is at least $1/2$. This success probability can be amplified to $1 - 2^{-m}$ by executing the algorithm m times and taking the union of all sets of solutions. 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 admits a *randomized FPTAS* [13], i. e., the algorithm succeeds in computing a $(1 - \varepsilon, 1 + \varepsilon)$ approximate Pareto curve with a probability of at least $1/2$. This randomized FPTAS yields also a randomized FPTAS for the multi-criteria cycle cover problem [11].

1.3 Previous Work

Most research on multi-criteria TSP is about heuristics for finding approximate solutions without any worst-case guarantee. We refer to Ehrgott and Gandibleux [6, 7] for a comprehensive survey.

The first result concerning computing approximate Pareto curves for the TSP is due to Angel et al. [1, 2], who considered Min-STSP restricted to edge weights 1 and 2. Ehrgott [5] considered a variant of MC-Min-STSP, where all objectives are encoded into a single objective by using some norm. MC-Min-STSP allows for a $(2+\varepsilon)$ approximation [11]. Bläser et al. [4] devised the first randomized approximations for MC-Max-STSP and MC-Max-ATSP. Their algorithms achieve ratios of $\frac{1}{k} - \varepsilon$ for k-Max-STSP and $\frac{1}{k+1} - \varepsilon$ for k-Max-ATSP. This has been improved to $2/3 - \varepsilon$ and $1/2 - \varepsilon$, respectively [10]. MC-Min-ATSP can be approximated with a factor of $\log_2 n + \varepsilon$ [10].

All approximation algorithms mentioned above deal with the special cases where we have either only min objectives or only max objectives. As far as we are aware, nothing is known so far about the approximability of multi-criteria TSP with both min and max objectives.

1.4 New Results

We present randomized approximation algorithms for k-Max-ℓ-Min-TSP for any $k, \ell \in \mathbb{N}$. As far as we are aware, this is the first paper that deals with approximation algorithms for multi-criteria TSP with both min and max objectives simultaneously. Our approximation algorithm for k-Max-ℓ-Min-ATSP computes $(1/2 - \varepsilon, \log_2 n + \varepsilon)$ approximate Pareto curves for any ℓ and k (Section 3). Our approximation algorithm for k-Max-ℓ-Min-STSP computes $(2/3 - \varepsilon, 4 + \varepsilon)$ approximate Pareto curves for any ℓ and k (Section 4). The running-times of our algorithms are polynomial in the input size for any fixed k, ℓ, and ε. But, since even the sizes of approximate Pareto curves are exponential in the number of objectives, it is unavoidable that the running-time is exponential in k and ℓ.

The main difficulty is that min and max objectives are different in nature: For max objectives, we have to collect as much weight as possible. If we have a substructure, i. e., a collection of paths that can be extended to a Hamiltonian cycle, then we can add any edges to actually get a Hamiltonian cycle. For min objectives, we have to be careful since adding any single heavy edge can deteriorate the approximation ratio.

The idea to deal with this difficulty is to first detect a collection of paths that have sufficient weight with respect to the max objectives. (In fact, we compute a set of collections of paths since a single collection does not suffice.) We will take care that these collections of paths are not too heavy with respect to the min objectives. After that, we connect our collections of paths to get Hamiltonian cycles. In this second step, we only pay attention to the min objectives; we already have enough weight with respect to the max objectives, and adding further edges does not decrease the weight.

In the next section, we introduce *decompositions*, which have already been used to approximate MC-Max-TSP [4,10]. Then we present our algorithms and their analyses in the subsequent sections. As a byproduct, our algorithms are simplified $1/2 - \varepsilon$ and $2/3 - \varepsilon$ approximation algorithms for MC-Max-ATSP and MC-Max-STSP, respectively. In particular, they avoid the recursion from k to $k - 1$ objectives that was used in the earlier approximation algorithms for MC-Max-TSP [4,10]. Finally, we consider some variants of the problem like combining asymmetric and symmetric objective functions (Section 5).

Due to space constraints, most proofs are omitted due to space constraints.

2 Decompositions

From now on, let $\eta_{k,\varepsilon} = \frac{\varepsilon^2}{2 \ln k} < 1$ for $\varepsilon > 0$. We assume $\varepsilon < \frac{1}{2 \ln k}$ throughout the paper. This is no restriction since the number k of max objectives is considered to be fixed. For $n \in \mathbb{N}$, let $[n] = \{1, 2, \ldots, n\}$.

We call a Hamiltonian cycle H a *ξ-heavy-weight Hamiltonian cycle* if there exists an $i \in [k]$ and an edge $e \in H$ such that $w_i(e) > \xi w(H)$. In this case, e is called a *ξ-heavy-weight edge* of H. If ξ is clear from the context, we also speak simply of a heavy-weight Hamiltonian cycle and a heavy-weight edge. Vice versa, H is a *ξ-light-weight Hamiltonian cycle* if it is not ξ-heavy-weight. Light-weight and heavy-weight cycle covers as well as heavy-weight edges of cycle covers are defined analogously.

A *decomposition* of a cycle cover C is a set $P \subseteq C$ of edges that consists solely of paths. The collection P of paths is obtained by removing a single edge of every cycle of C. The set P is called a γ decomposition if $w(P) \geq \gamma w(C)$. Decompositions play a crucial role in approximating MC-Max-TSP: We can add edges to a collection P of paths to get a Hamiltonian cycle. Thus, if C allows for an γ decomposition P, then we can find a Hamiltonian cycle $H \supseteq P$ with $w(H) \geq$

$w(P) \geq \gamma w(C)$. For our algorithms, we exploit that $(1/2 - \varepsilon)$ decompositions of directed $\eta_{k,\varepsilon}$-light-weight cycle covers and $(2/3-\varepsilon)$ decompositions of undirected $\eta_{k,\varepsilon}$-light-weight cycle covers exist and can be found in polynomial time [10].

We call the procedure that finds decompositions Decompose with parameters C, w, and ε: C is a cycle cover (directed or undirected), $w = (w_1, \ldots, w_k)$ are k edge weights, and $\varepsilon > 0$. Then Decompose(C, w, ε) returns a $(1/2 - \varepsilon)$- or $(2/3 - \varepsilon)$-decomposition $P \subseteq C$, depending on whether C is directed or undirected, provided that C is an $\eta_{k,\varepsilon}$-light-weight cycle cover.

In addition to Decompose, we use the following existing algorithms for our algorithms: CyCo-Approx denotes the randomized FPTAS for cycle covers: on input $(G, w, d, \varepsilon, p)$, CyCo-Approx returns a $(1 - \varepsilon, 1 + \varepsilon)$ approximate Pareto curves of cycle covers with respect to (G, w, d) with a success probability of at least $1 - p$. We use CyCo-Approx for computing both undirected and directed cycle covers. If either d or w is missing, CyCo-Approx computes a $(1 - \varepsilon)$ or $(1 + \varepsilon)$ approximate Pareto curve with respect to w or d, respectively.

MST-Approx denotes the deterministic FPTAS for multi-criteria spanning trees of minimum weight [13]: MST-Approx(G, d, ε) computes a $(1+\varepsilon)$ approximate Pareto curve of spanning trees of the instance (G, d).

By MinATSP-Approx, we denote the $(\log_2 n + \varepsilon)$ approximation algorithm for MC-Min-ATSP [10]: On input (G, d, ε, p), MinATSP-Approx computes a $(\log_2 n + \varepsilon)$ approximate Pareto curve for MC-Min-ATSP for the instance (G, d) with a success probability of at least $1 - p$.

3 Asymmetric Multi-Criteria TSP

3.1 Preparation for MC-ATSP

In this section, we focus our attention on the max objectives w. For a graph $G = (V, E)$ and a subset $K \subseteq E$ of G's edges, we obtain G_{-K} by contracting all edges of K. Contracting a single edge (u, v) means removing all outgoing edges of u, removing all incoming edges of v, and identifying u and v. (The set K will always be such that no conflicts arise during contraction. In particular, the order in which the edges are contracted does not matter.) Analogously, for a Hamiltonian cycle H and edges K, we obtain a Hamiltonian cycle H_{-K} of G_{-K} by contracting the edges in K. We will usually have $K \subseteq H$ in this case.

If (G, w, d) is an instance for a multi-criteria TSP problem, then (G_{-K}, w, d) denotes the instance with w and d modified according to the edge contractions.

For any Hamiltonian cycle H, let $\zeta_i = \max\{w_i(e) \mid e \in H\}$ be the weight of the heaviest edge with respect to the i-th objective. Let $\zeta = \zeta(H) = (\zeta_1, \ldots, \zeta_k)$. We will distinguish between H being a light-weight cycle cover, i. e., all components of $\zeta = \zeta(H)$ are small, and H being a heavy-weight cycle cover, i. e., there is some i such that ζ_i is large. From the edge weights w and ζ, we obtain new edge weights w^ζ by setting the weight of all edges that are heavier than any edge in H to 0:

$$w^\zeta(e) = \begin{cases} w(e) & \text{if } w(e) \leq \zeta \text{ and} \\ 0 & \text{if } w_i(e) > \zeta_i \text{ for some } i. \end{cases}$$

This does not affect the weight of H since all edges $e \in H$ fulfil $w(e) \leq \zeta$. The reason for this definition is the following: Assume that H is a light-weight cycle cover, and assume that we have a $(1-\varepsilon)$ approximate Pareto curve $\mathcal{C}^{\zeta(H)}$ of cycle covers with respect to $w^{\zeta(H)}$. Then $\mathcal{C}^{\zeta(H)}$ contains a light-weight cycle cover whose weight is close to H's weight. This is stated more precisely in the following lemma.

Lemma 1 (Manthey [10]). *Let $\varepsilon > 0$ be sufficiently small. Let H be an $(\eta_{k,\varepsilon/2} - (\frac{\varepsilon}{2})^3)$-light-weight Hamiltonian cycle. Let $\zeta = \zeta(H)$, and let $\mathcal{C}^\zeta$ be a $(1-\frac{\varepsilon}{2})$ approximate Pareto curve of cycle covers with respect to w^ζ.*

Then $\mathcal{C}^\zeta$ contains a cycle cover C with $w^\zeta(C) \geq (1-\frac{\varepsilon}{2}) \cdot w(H)$ and $w^\zeta(e) \leq \eta_{k,\varepsilon/2} \cdot w^\zeta(C)$ for all $e \in C$. This cycle cover C yields a decomposition $P \subseteq C$ with $w(P) \geq (1/2-\varepsilon) \cdot w(H)$.

This is all we need so far for dealing with light-weight Hamiltonian cycles. Next, we deal with heavy-weight Hamiltonian cycles $\tilde{H}$. Of course it can happen that we somehow get a light-weight cycle cover C that approximates $\tilde{H}$, i. e., $w(C) \geq (1-\varepsilon)w(\tilde{H})$. In this case, we can apply decomposition and are done.

However, we cannot guarantee that we find such a cycle cover, not even that such a cycle cover exists. Thus, heavy-weight Hamiltonian cycles need special treatment. The idea how to deal with them is to collect a few number of heavy-weight edges. This should be done such that the following properties are met: The collection should contain a $1/2-\varepsilon$ fraction of the weight of $\tilde{H}$ with respect to some objective functions. And the rest of $\tilde{H}$, after all edges of the collection have been contracted, should be a light-weight Hamiltonian cycle. This would allow us to use decomposition for the rest of $\tilde{H}$. Let $f(k,\varepsilon) = k \cdot \lceil \frac{\log(\frac{1}{2}+\varepsilon)}{\log(1-\eta_{k,\varepsilon/2}+(\frac{\varepsilon}{2})^3)} \rceil$. Our goal is now to prove that for every Hamiltonian cycle H, there exists a set $K \subseteq H$ of cardinality at most $f(k,\varepsilon)$ such that H_{-K} is a $\big(\eta_{k,\varepsilon/2} - (\frac{\varepsilon}{2})^3\big)$-light-weight Hamiltonian cycle.

Lemma 2. *For every H and every $\varepsilon > 0$, there exists a subset $K \subseteq H$ such that $|K| \leq f(k,\varepsilon)$ and, for every $i \in [k]$, we have*

1. *$w_i(K) \geq (1/2-\varepsilon)w_i(H)$ or*
2. *$w_i(e) \leq \big(\eta_{k,\varepsilon/2} - (\frac{\varepsilon}{2})^3\big)w_i(H_{-K})$ for all $e \in H_{-K}$.*

3.2 Approximation Algorithm for MC-ATSP

From the results of the previous section, we know that ζ and K exist such that, for every $\tilde{H}$, we will find an appropriate light-weight cycle cover that eventually yields a tour H whose weight approximates $\tilde{H}$'s weight. To actually obtain an algorithm, we have to find K and ζ. But there is only a polynomial number of possibilities for ζ and K: For all ζ and for all $i \in [k]$, we can assume that there is an edge with $w_i(e) = \zeta_i$. Thus, there are at most $O(n^2)$ choices for ζ_i, hence at most $O(n^{2k})$ in total. The cardinality of K is bounded in terms of $f(k,\varepsilon)$ as we have shown in the lemma above. For fixed k and ε, there is only

$\mathcal{P}_{\text{TSP}} \leftarrow$ ATSP-Approx(G, w, d, ε)
input: directed complete graph $G = (V, E)$, $w : E \to \mathbb{Q}_+^k$, $d : E \to \mathbb{Q}_+^\ell$, $\varepsilon > 0$
output: $(1/2 - \varepsilon, \log_2 n + \varepsilon)$ approximate Pareto curve $\mathcal{P}_{\text{TSP}}$ for k-Max-ℓ-Min-ATSP with a success probability of at least $1/2$
for all path covers $K \subseteq E$ with $|K| \le f(k, \varepsilon)$ and bounds ζ **do**
  $\mathcal{C}_{K,\zeta} \leftarrow$ CyCo-Approx$\big(G_{-K}, w^\zeta, d, \frac{\varepsilon}{2}, \frac{1}{4n^{2k+f(k,\varepsilon)}}\big)$
  for all $I \subseteq [k]$ and $C \in \mathcal{C}_{K,\zeta}$ **do**
    if $w_I^\zeta(e) \le \eta_{k,\varepsilon/2} \cdot w_I^\zeta(C)$ for all $e \in C$ **then**
      $P \leftarrow$ Decompose$\big(C, w_I^\zeta, \frac{\varepsilon}{2}\big)$
      let V' be the start-points of paths of P
      $\mathcal{P}'_{\text{TSP}} \leftarrow$ MinATSP-Approx$\big(V', d, \frac{\varepsilon}{2}, \frac{1}{2^k 4n^{2k+f(k,\varepsilon)}|\mathcal{C}_{K,\zeta}|}\big)$
      for all $H' \in \mathcal{P}'_{\text{TSP}}$ **do**
        $A \leftarrow H' \cup C$
        obtain a tour H'' from the Eulerian set A of edges with $H'' \supseteq P$
        combine H'' and K to a tour H; add H to $\mathcal{P}_{\text{TSP}}$

Algorithm 1: ATSP-Approx: Approximation algorithm for MC-ATSP.

a polynomial number of subsets of cardinality at most $f(k, \varepsilon)$. We can restrict K to be a path cover, which is an acyclic set of edges such that both indegree and outdegree of each vertex is at most one. Of course, the running-time of our algorithm is exponential in the number k of max objectives. But this is unavoidable since the sizes of approximate Pareto curves can be exponential in the number of objectives. In the following, w_I for a set $I \subseteq [k]$ denotes the vector of edge weights restricted to the components in I. Instead of taking edges one-by-one as in the proof of Lemma 2, we take all edges at once. This means that we take a subset of the edges of cardinality at most $f(k, \varepsilon)$. Furthermore, we do not distinguish between light-weight and heavy-weight Hamiltonian cycles: light-weight Hamiltonian cycles are simply those for which $K = \emptyset$ works.

The min objectives remain to be taken into account. The main idea behind the algorithm is first to collect enough weight with respect to the max objectives. This gives us a collection of paths that fulfil the weight requirements for the max objectives. We have to be careful not to get too much weight with respect to the min objectives. After that we connect the paths using MinATSP-Approx, which is the approximation algorithm for MC-Min-ATSP. Overall, we get ATSP-Approx (Algorithm 1) and the following theorem.

Theorem 1. *For every* $\varepsilon > 0$, ATSP-Approx *(Algorithm 1) is a randomized* $(1/2 - \varepsilon, \log_2 n + \varepsilon)$ *approximation algorithm for k-Max-ℓ-Min-ATSP for any* $k, \ell \in \mathbb{N}$. *For fixed* ε, k, *and* ℓ, *its running-time is polynomial in the input size.*

Proof. For want of space, we only analyze the approximation ratio. To do this, we assume that all randomized computations are successful. We have to show that for every Hamiltonian cycle $\tilde{H}$, there exists a Hamiltonian cycle $H \in \mathcal{P}_{\text{TSP}}$ with $w(H) \ge (1/2 - \varepsilon) w(\tilde{H})$ and $d(H) \le (\log_2 n + \varepsilon) d(\tilde{H})$. Thus, let $\tilde{H}$ be an arbitrary Hamiltonian cycle. By Lemma 2, there exists a set $K \subseteq \tilde{H}$ of cardinality at most $f(k, \varepsilon)$ and a set $I \subseteq [k]$ with the following properties:

– For every $i \in [k] \setminus I$, we have $w_i(K) \geq (1/2 - \varepsilon) w_i(\tilde{H})$.
– For every $i \in I$ and for every edge $e \in H_{-K}$, we have $w_i(e) \leq \big(\eta_{k,\varepsilon/2} - (\frac{\varepsilon}{2})^3\big) w_i(H_{-K})$.

Let $\zeta = \zeta(H_{-K})$. According to Lemma 1, the set $\mathcal{C}_{K,\zeta}$ contains a cycle cover C with the following properties:

– C is a $\eta_{k,\varepsilon/2}$-light-weight cycle cover with respect to w_I.
– $w_i(C) \geq (1 - \frac{\varepsilon}{2}) w_i(\tilde{H})$ for every $i \in I$.
– $d(C) \leq (1 + \frac{\varepsilon}{2}) d(\tilde{H}_{-K})$.

This means that there exists a decomposition $P \subseteq C$ such that $w_i(P) \geq (1/2 - \varepsilon) w_i(\tilde{H}_{-K})$ and P consists of at most $n/2$ paths. The former follows from Lemma 1. The latter holds since H_{-K} has at most n vertices and every cycle of C consists of at least two vertices, which implies that every connected component of P consists of at least two vertices.

Thus, V' has at most $n/2$ vertices. This implies that $\mathcal{P}'_{\text{TSP}}$ contains a Hamiltonian cycle H' with $d(H') \leq \big(\log_2(\frac{n}{2}) + \frac{\varepsilon}{2}\big) d(H_{-K})$.

We obtain the Hamiltonian cycle H'' of $V \setminus K$ 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 $d(v, x) \leq d(v, u) + d(u, x)$. This yields

$$\begin{aligned} d(H'') &\leq d(C) + d(H') \leq d(C) + \left(\log_2\left(\tfrac{n}{2}\right) + \tfrac{\varepsilon}{2}\right) \cdot d(H_{-K}) \\ &\leq \left(1 + \tfrac{\varepsilon}{2} + \log_2\left(\tfrac{n}{2}\right) + \tfrac{\varepsilon}{2}\right) \cdot d(H_{-K}) = (\log_2 n + \varepsilon) \cdot d(H_{-K}). \end{aligned}$$

We observe that $d(\tilde{H}) = d(K) + d(\tilde{H}_{-K})$ since $K \subseteq \tilde{H}$. Furthermore, $d(H) = d(H'') + d(K)$. Thus, $d(H) \leq (\log_2 n + \varepsilon) \cdot d(H_{-K}) + d(K) \leq (\log_2 n + \varepsilon) \cdot d(\tilde{H})$. In addition, we have $w_i(H) \geq w_i(K) \geq \left(\frac{1}{2} - \varepsilon\right) \cdot w(\tilde{H})$ for every $i \in I$ and $w_i(H) \geq w_i(P) + w_i(K) \geq \left(\frac{1}{2} - \varepsilon\right) w_i(\tilde{H}_{-K}) + w_i(K) \geq \left(\frac{1}{2} - \varepsilon\right) w_i(\tilde{H})$ for every $i \in [k] \setminus I$. This proves the approximation ratio. □

4 Symmetric Multi-Criteria TSP

Of course, ATSP-Approx works also for MC-STSP. However, this ignores d and w being symmetric. In this section, we present a $(2/3 - \varepsilon, 4 + \varepsilon)$ approximation algorithm for MC-STSP.

4.1 Preparation for MC-STSP

As we did for ATSP, we first focus our attention on the max objectives w.

For our approximation algorithm, we need counterparts of Lemmas 1 and 2. The following function g plays a similar role as f in the directed case: $g(k, \varepsilon) = k \cdot \left\lceil \frac{\log(\frac{1}{3} + \frac{\varepsilon}{3})}{\log(1 - \eta_{k,\varepsilon/3} + (\frac{\varepsilon}{3})^3)} \right\rceil$. For bounds $\zeta = (\zeta_1, \ldots, \zeta_k) \in \mathbb{Q}_+^k$, we define w^ζ in the same way as for directed graphs. We have the following counterpart of Lemma 1.

Lemma 3 (Manthey [10]). *Let $\varepsilon > 0$ be arbitrary. Let $\tilde{H}$ be an undirected Hamiltonian cycle that is $\bigl(\eta_{k,\varepsilon/3} - (\frac{\varepsilon}{3})^3\bigr)$-light. Let $\zeta = \zeta(\tilde{H})$, and let $\mathcal{C}^\zeta$ be a $\bigl(1 - \frac{\varepsilon}{3}\bigr)$ approximate Pareto curve of cycle covers with respect to w^ζ.*

Then $\mathcal{C}^\zeta$ contains a cycle cover C with $w^\zeta(C) \geq \bigl(1 - \frac{\varepsilon}{3}\bigr) w(\tilde{H})$ and $w^\zeta(e) \leq \eta_{k,\varepsilon/3} w^\zeta(C)$ for all $e \in C$. This cycle cover C yields a decomposition $P \subseteq C$ with $w(P) \geq \bigl(\frac{2}{3} - \frac{2\varepsilon}{3}\bigr) w(\tilde{H})$.

However, the main difficulty when dealing with STSP is that contractions are no longer possible. If we contract an edge in the straight-forward way, we obtain a directed instance. Since we aim at better approximation ratios for STSP than we have for ATSP, something more sophisticated has to be done. Instead of taking only single edges, we take longer paths. We do not contract these paths, but set the weights of all edges incident to vertices on the path to 0. In this way, we can later remove the edges of a cycle that traverse these vertices, and then we can add the edges of the path. The problem is that we might lose the two edges at the ends of the paths; we cannot force them to be in a cycle cover in the same way. However, as the following lemma shows, we can choose the two edges at the end such that they contribute only little to the weight of the Hamiltonian cycle. Thus, we do not lose too much weight and are still able to achieve a good approximation ratio. The following lemma shows that any sufficiently long path contains an edge that is light with respect to all objectives $w_1, \ldots, w_k$.

Lemma 4 (Manthey [10]). *Let $\tilde{H}$ be a Hamiltonian cycle on n vertices, and let $e_1, \ldots, e_m$ be any m distinct edges of $\tilde{H}$. Then there exists a $z \in [m]$ such that $w(e_z) \leq \frac{k}{m} \cdot w(\tilde{H})$.*

Now let $\tilde{H}$ be a Hamiltonian cycle, and let $K \subseteq \tilde{H}$. Let $L = L(K) = \{v \in V \mid \exists e \in K : v \in e\}$ be the set of vertices incident to edges in K. Let w^{-L} be defined by $w^{-L}(e) = w(e)$ if $e \cap L = \emptyset$ and $w^{-L}(e) = 0$ if $e \cap L \neq \emptyset$. This means that the weight of edges incident to L is set to 0, which includes the edges in K. But there are more edges whose weight is affected by w^{-L}: Let

$$T = T(K) = \{e \in \tilde{H} \mid e \notin K, e \cap L(K) \neq \emptyset\}$$

be the set of edges that have exactly one endpoint in L. The weights of these edges are set to 0 in w^{-L}, but we cannot force them to be in any cycle cover as mentioned above. (They are the edges at the ends of the paths in K.) The following lemma is the undirected counterpart of Lemma 2. In particular, it takes care of the set T. This set T is only needed for the analysis and not for the algorithm.

Lemma 5. *For every Hamiltonian cycle H and every $\varepsilon > 0$, there exists a subset $K \subseteq H$ of at most $g(k, \varepsilon)$ paths, each of length at most $\frac{6k}{\varepsilon} g(k, \varepsilon)$ with the following properties: Let $L = L(K)$ and $T = T(K)$. For every $i \in [k]$, we have*

1. *$w_i(K) \geq (2/3 - \varepsilon) w_i(H)$ or*
2. *$w_i^{-L}(e) \leq \bigl(\eta_{k,\varepsilon/3} - (\frac{\varepsilon}{3})^3\bigr) w_i^{-L}(H)$ for all $e \in H$.*

Furthermore, we have $w(T) \leq \frac{2\varepsilon}{3} w(H)$.

$\mathcal{P}_{\text{TSP}} \leftarrow$ STSP-Approx(G, w, d, ε)

input: undirected complete graph $G = (V, E)$, $w : E \to \mathbb{Q}_+^k$, $d : E \to \mathbb{Q}_+^\ell$, $\varepsilon > 0$

output: $(2/3 - \varepsilon, 4 + \varepsilon)$ approximate Pareto curve $\mathcal{P}_{\text{TSP}}$ for k-Max-ℓ-Min-ATSP with a success probability of at least 1/2

for all $K \subseteq E$ consisting of $\leq g(k, \varepsilon)$ paths of length $\leq 6kg(k, \varepsilon)$ and all ζ **do**
    $L \leftarrow L(K)$
    $\mathcal{C}_{L,\zeta} \leftarrow$ CyCo-Approx$\big(G, w^{-L,\zeta}, \frac{\varepsilon}{4}, \frac{1}{n^{2k+6kg^2(k,\varepsilon)}}\big)$
    for all $I \subseteq [k]$ and $C \in \mathcal{C}_{L,\zeta}$ **do**
        if $w_I^{-L,\zeta}(e) \leq \eta_{k,\varepsilon/3} \cdot w_I^{-L,\zeta}(C)$ for all $e \in C$ **then**
            $P \leftarrow$ Decompose$\big(C, w_I^{-L,\zeta}, \frac{\varepsilon}{4}\big)$
            remove edges of weight 0 from P
            choose one end-point of each path of P and K to obtain V'
            let G' be the corresponding graph
            $\mathcal{T} \leftarrow$ MST-Approx$\big(G', d, \frac{\varepsilon}{4}\big)$
            for all $T \in \mathcal{T}$ **do**
                combine T, P, and K to a spanning tree T' of G
                duplicate each edge of T' to get an Eulerian graph T''
                traverse T'', take shortcuts to get a tour $H \supseteq P \cup K$; add H to $\mathcal{P}_{\text{TSP}}$

Algorithm 2: STSP-Approx: Approximation algorithm for MC-STSP.

4.2 Approximation Algorithm for MC-STSP

The main difficulty in getting approximation ratios for MC-STSP is threefold: First, we have to be more careful than for MC-ATSP since contractions are impossible. When inserting the edges of the set K, we have to take into account two points: First, we need all edges of K since we need the weight for the max objectives. Second, we cannot afford to add arbitrary edges to build a Hamiltonian cycle since this might add too much weight with respect to the min objectives. Third, concerning the approximation ratio, we will construct Eulerian graphs from which we obtain the Hamiltonian cycles by taking shortcuts. However, on the one hand, we have to make sure that none of the edges of K is removed by taking shortcuts. On the other hand, this gives us another factor of 2 in the approximation ratio. (The problem is that Christofides' algorithm for MC-Min-STSP gives us only a ratio of $2+\varepsilon$ instead of 3/2 as it does for Min-STSP with a single objective.) We deal with these issues in the proof of the main theorem of this section. Overall, we obtain Algorithm 2 (STSP-Approx) and the following result.

Theorem 2. *For every $\varepsilon > 0$,* STSP-Approx *(Algorithm 2) is a randomized $(2/3-\varepsilon, 4+\varepsilon)$ approximation algorithm for k-Max-ℓ-Min-STSP for any $k, \ell \in \mathbb{N}$. For fixed ε, k, and ℓ, its running-time is polynomial in the input size.*

5 Variants

Since k and ℓ are usually quite small, a natural question is if the approximation ratios can be improved for particular values of k and ℓ.

Our first observation is that k-Max-1-Min-STSP allows for a $(2/3-\varepsilon, 3.5+\varepsilon)$ approximation: Instead of using the spanning tree heuristic in lines 10 to 13, we use Christofides' algorithm [14]. More general and along the same lines: If ℓ-Min-STSP can be approximated with a ratio of s_ℓ, then this yields a $(2/3-\varepsilon, s_\ell+2+\varepsilon)$ approximation algorithm for k-Max-ℓ-Min-STSP.

Our second observation concerns ATSP: If ℓ-Min-ATSP can be approximated with a ratio of $s_\ell(n)$ on graphs with n vertices, then this yields a $(1/2-\varepsilon, 1+s_\ell(n/2)+\varepsilon)$ approximation algorithm for k-Max-ℓ-Min-ATSP. This follows immediately from the analysis in Section 3. In particular, for $\ell = 1$, we obtain a $(1/2-\varepsilon, \frac{2}{3}\log_2 n + \frac{1}{3} + \varepsilon)$ approximation using the algorithm of Feige and Singh for Min-ATSP [8].

Finally, an obvious variant of multi-criteria TSP that has not been analyzed yet is a combination of ATSP and STSP: Some objectives are asymmetric, while others are symmetric. The difficulty with this variant is that, for asymmetric objectives, only cycle covers with a minimum cycle length of two can be computed efficiently. Thus, if also the symmetric objectives require cycle cover computations, which is the case for symmetric max objectives, it seems hard to get approximation ratios better than the trivial ratios that we obtain by using the ATSP algorithms for both symmetric and asymmetric objectives.

One setting, however, allows for better ratios: If the max objectives are asymmetric and the min objectives are symmetric, then a straightforward combination of ATSP-Approx (Algorithm 1) and STSP-Approx (Algorithm 2) gives a ratio of $(1/2-\varepsilon, 4+\varepsilon)$: We run ATSP-Approx until we have enough weight with respect to the max objectives w. Then we switch to STSP-Approx to connect the components. We do not lose any weight with respect to w by connecting the components, although the max objectives w are asymmetric.

6 Concluding Remarks

We have presented approximation algorithms for multi-criteria traveling salesman problems that have min and max objectives simultaneously. Our algorithms work for any fixed number of minimization and maximization objectives. They are randomized and have polynomial running-time. The approximation ratios obtained, $(1/2-\varepsilon, \log_2 n+\varepsilon)$ for MC-ATSP and $(2/3-\varepsilon, 4+\varepsilon)$ for MC-STSP, match the approximation ratios for multi-criteria TSP with only maximization or only minimization problems, except for the Min-STSP part of MC-STSP. For this, the ratio is only $4+\varepsilon$, compared to $2+\varepsilon$ for MC-Min-STSP. This raises the questions whether this $4+\varepsilon$ can be improved. More precisely: If there exists an r_ℓ approximation algorithm for ℓ-Min-STSP, does this yield a $(2/3-\varepsilon, r_\ell)$ approximation algorithm for k-Max-ℓ-Min-STSP? So far, we only get a performance ratio of $(2/3-\varepsilon, r_\ell+2+\varepsilon)$ according to Section 5.

To simplify the analysis of the approximability of multi-criteria TSP, it would be nice if any improvement for k-Max-STSP also yields an improvement for k-Max-ℓ-Min-STSP: Assume that k-Max-STSP can be approximated with a ratio of s_k and ℓ-Min-STSP can be approximated with a ratio of r_ℓ. Does this yield a (r_k, s_ℓ) approximation for k-Max-ℓ-Min-STSP? Or at least a (f_k, g_ℓ) approximation for some non-trivial functions f_k and g_ℓ that depend on r_k and s_ℓ? The same question arises for k-Max-ATSP, ℓ-Min-ATSP, and k-Max-ℓ-Min-ATSP.

Finally, we ask whether there are also faster and deterministic algorithms for multi-criteria TSP. The algorithms presented here use randomness only because no deterministic FPTAS for multi-criteria cycle covers is known. Maybe either the randomized FPTAS can be derandomized or cycle covers as an intermediate step can be avoided at all.

References

1. Eric Angel, Evripidis Bampis, and Laurent Gourvés. Approximating the Pareto curve with local search for the bicriteria TSP(1,2) problem. *Theoret. Comput. Sci.*, 310(1–3):135–146, 2004.
2. Eric Angel, Evripidis Bampis, Laurent Gourvès, and Jérǫme Monnot. (Non-)approximability for the multi-criteria TSP(1,2). In *Proc. 15th Int. Symp. on Fundamentals of Computation Theory (FCT)*, vol. 3623 of *Lecture Notes in Comput. Sci.*, pp. 329–340. Springer, 2005.
3. Markus Bläser and Bodo Manthey. Approximating maximum weight cycle covers in directed graphs with weights zero and one. *Algorithmica*, 42(2):121–139, 2005.
4. Markus Bläser, Bodo Manthey, and Oliver Putz. Approximating multi-criteria Max-TSP. In *Proc. 16th Ann. European Symp. on Algorithms (ESA)*, vol. 5193 of *Lecture Notes in Comput. Sci.*, pp. 185–197. Springer, 2008.
5. Matthias Ehrgott. Approximation algorithms for combinatorial multicriteria optimization problems. *Int. Trans. Oper. Res.*, 7(1):5–31, 2000.
6. Matthias Ehrgott. *Multicriteria Optimization.* Springer, 2005.
7. Matthias Ehrgott and Xavier Gandibleux. A survey and annotated bibliography of multiobjective combinatorial optimization. *OR Spectrum*, 22(4):425–460, 2000.
8. Uriel Feige and Mohit Singh. Improved approximation ratios for traveling salesperson tours and paths in directed graphs. In *Proc. 10th Int. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, vol. 4627 of *Lecture Notes in Comput. Sci.*, pp. 104–118. Springer, 2007.
9. Haim Kaplan, Moshe Lewenstein, Nira Shafrir, and Maxim I. Sviridenko. Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. *J. ACM*, 52(4):602–626, 2005.
10. Bodo Manthey. On approximating multi-criteria TSP. In *Proc. 26th Int. Symp. on Theoretical Aspects of Computer Science (STACS)*, pp. 637–648, 2009.
11. Bodo Manthey and L. Shankar Ram. Approximation algorithms for multi-criteria traveling salesman problems. *Algorithmica*, 53(1):69–88, 2009.
12. Katarzyna Paluch, Marcin Mucha, and Aleksander Madry. A 7/9 approximation algorithm for the maximum traveling salesman problem. In *Proc. 12th Int. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, vol. 5687 of *Lecture Notes in Comput. Sci.*, pp. 298–311. Springer, 2009.
13. Christos H. Papadimitriou and Mihalis Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *Proc. 41st Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pp. 86–92. IEEE, 2000.
14. Vijay V. Vazirani. *Approximation Algorithms.* Springer, 2001.