Symposium on Theoretical Aspects of Computer Science 2009 (Freiburg), pp. 637–648
www.stacs-conf.org

ON APPROXIMATING MULTI-CRITERIA TSP

BODO MANTHEY [1]

[1] Saarland University, Computer Science, Postfach 151150, 66041 Saarbrücken, Germany
E-mail address: manthey@cs.uni-sb.de

Abstract. We present approximation algorithms for almost all variants of the multi-criteria traveling salesman problem (TSP), whose performances are independent of the number k of criteria and come close to the approximation ratios obtained for TSP with a single objective function.

We present randomized approximation algorithms for multi-criteria maximum traveling salesman problems (Max-TSP). For multi-criteria Max-STSP, where the edge weights have to be symmetric, we devise an algorithm that achieves an approximation ratio of $2/3 - \varepsilon$. For multi-criteria Max-ATSP, where the edge weights may be asymmetric, we present an algorithm with an approximation ratio of $1/2 - \varepsilon$. Our algorithms work for any fixed number k of objectives. To get these ratios, we introduce a decomposition technique for cycle covers. These decompositions are optimal in the sense that no decomposition can always yield more than a fraction of 2/3 and 1/2, respectively, of the weight of a cycle cover. Furthermore, we present a deterministic algorithm for bi-criteria Max-STSP that achieves an approximation ratio of $61/243 \approx 1/4$.

Finally, we present a randomized approximation algorithm for the asymmetric multi-criteria minimum TSP with triangle inequality (Min-ATSP). This algorithm achieves a ratio of $\log n + \varepsilon$. For this variant of multi-criteria TSP, this is the first approximation algorithm we are aware of. If the distances fulfil the γ-triangle inequality, its ratio is $1/(1-\gamma) + \varepsilon$.

1. Multi-Criteria Traveling Salesman Problem

Traveling Salesman Problem. The traveling salesman problem (TSP) is one of the most famous combinatorial optimization problems. Given a graph, the goal is to find a Hamiltonian cycle (also called a *tour*) of maximum or minimum weight (Max-TSP or Min-TSP). 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 of maximum weight. The weight of a Hamiltonian cycle (or, more general, of any set of edges) is the sum of the weights of its edges. If G is undirected, we have Max-STSP (symmetric TSP). If G is directed, we obtain Max-ATSP (asymmetric TSP). An instance of Min-TSP is also a complete graph G with edge weights w that fulfil the triangle inequality: $w(u, v) \leq w(u, x) + w(x, v)$ for all $u, v, x \in V$. The goal is to find a tour of minimum weight. We have Min-STSP if G is undirected and Min-ATSP if G is directed. In this paper, we only consider the latter. If we restrict the instances to fulfil the γ-triangle inequality $(w(u, v) \leq \gamma \cdot (w(u, x) + w(x, v))$ for all distinct $u, v, x \in V$

Key words and phrases: Approximation algorithms, traveling salesman, multi-criteria optimization.

© B. Manthey
Creative Commons Attribution-NoDerivs License

and $\gamma \in [\frac{1}{2}, 1)$), then we obtain Min-$\gamma$-ATSP. All variants introduced are NP-hard and APX-hard (Min-γ-ATSP is hard for $\gamma > \frac{1}{2}$). Thus, we have to content ourselves with approximate solutions. The currently best approximation algorithms for Max-STSP and Max-ATSP achieve approximation ratios of 61/81 [7] and 2/3 [14]. 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 [11]. Min-γ-ATSP allows for an approximation ratio of $\min\{\frac{\gamma}{1-\gamma}, \frac{1+\gamma}{2-\gamma-\gamma^3}\}$ [5, 6].

Cycle covers are often used for designing approximation algorithms for the TSP [5, 14, 11, 6, 15, 7]. A cycle cover is a set of vertex-disjoint cycles such that every vertex is part of exactly one cycle. The idea is to compute an initial cycle cover and then to join the cycles to obtain a Hamiltonian cycle. This is called *subtour patching* [13]. Hamiltonian cycles are special cases of cycle covers that consist of a single cycle. Thus, the weight of a maximum-weight cycle cover bounds the weight of a maximum-weight Hamiltonian cycle from above, and the weight of a minimum-weight cycle cover is a lower bound for the weight of a minimum-weight Hamiltonian cycle. In contrast to Hamiltonian cycles, cycle covers of optimal weight can be computed efficiently by reduction to matching problems [1].

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 maximizing, e.g., the number of sights along the way. This leads to k-criteria variants of the TSP (k-C-Max-STSP, k-C-Max-ATSP, k-C-Min-STSP, k-C-Min-ATSP for short; if the number of criteria does not matter, we will also speak of MC-Max-STSP etc.). With respect to a single criterion, the term "optimal solution" is well-defined. However, if several criteria are involved, there is no natural notion of a best choice, and we have to be content with trade-off solutions. The goal of multi-criteria optimization is to cope with this dilemma. To transfer the concept of optimal solutions to multi-criteria problems, the notion of *Pareto curves* was introduced (cf. Ehrgott [9]). A Pareto curve is a set of solutions that can be considered optimal.

We introduce the following terms only for maximization problems. After that, we briefly state the differences for minimization problems. An instance of k-C-Max-TSP is a complete graph G with edge weights $w_1, \ldots, w_k : E \to \mathbb{Q}_+$. A tour H *dominates* another tour H' if $w_i(H) \geq w_i(H')$ for all $i \in [k] = \{1, \ldots, k\}$ and $w_i(H) > w_i(H')$ for at least one i. 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 maximization problems, k-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 often NP-hard to compute even for otherwise easy optimization problems. Third, the TSP is NP-hard already with one objective function, and optimization problems do not become easier with more objectives involved. Therefore, we have to be satisfied with approximate Pareto curves.

For simpler notation, let $w(H) = (w_1(H), \ldots, w_k(H))$. Inequalities are meant component-wise. A set $\mathcal{P}$ of Hamiltonian cycles of V is called an α *approximate Pareto curve* for (G, w) if the following holds: For every tour H', there exists a tour $H \in \mathcal{P}$ with $w(H) \geq \alpha w(H')$. We have $\alpha \leq 1$, and a 1 approximate Pareto curve is a Pareto curve. (This is not precisely true if there are several solutions whose objective values agree. But this is inconsequential here, and we will not elaborate on it for the sake of clarity.)

An algorithm is called an α approximation algorithm if, given G and w, it computes an α approximate Pareto curve. It is called a randomized α 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. (We can also remove solutions from this union that are dominated by other solutions in the union, but this is not required by the definition of an approximate Pareto curve.) Again, the concepts can be transfered easily to other maximization problems.

Papadimitriou and Yannakakis [18] showed that $(1-\varepsilon)$ approximate Pareto curves of size polynomial in the instance size and $1/\varepsilon$ exist. The technical requirement for the existence is that the objective values of all solutions for an instance X are bounded from above by $2^{p(N)}$ for some polynomial p, where N is the size of X. This is fulfilled in most optimization problems and in particular in our case. However, they only prove the existence, and for many optimization problems it is unclear how to actually find an approximate Pareto curve.

A *fully polynomial time approximation scheme* (FPTAS) for a multi-criteria optimization problem computes $(1-\varepsilon)$ approximate Pareto curves in time polynomial in the size of the instance and $1/\varepsilon$ for all $\varepsilon > 0$. Multi-criteria maximum-weight matching admits a *randomized FPTAS* [18], i.e., the algorithm succeeds in computing a $(1-\varepsilon)$ approximate Pareto curve with constant probability. This randomized FPTAS yields also a randomized FPTAS for the multi-criteria maximum-weight cycle cover problem [17].

To define Pareto curves and approximate Pareto curves also for minimization problems, in particular for MC-Min-STSP and MC-Min-ATSP, we have to replace all "$\geq$" and "$>$" above by "$\leq$" and "$<$". Furthermore, α approximate Pareto curves are now defined for $\alpha \geq 1$, and an FPTAS has to achieve an approximation ratio of $1+\varepsilon$. There also exists a randomized FPTAS for the multi-criteria minimum-weight cycle cover problem.

Related Work. For an overview of the literature about multi-criteria optimization, including multi-criteria TSP, we refer to Ehrgott and Gandibleux [10]. Angel et al. [2, 3] considered Min-STSP restricted to edge weights 1 and 2. They analyzed a local search heuristic and proved that it achieves an approximation ratio of 3/2 for $k=2$ and of $\frac{2k}{k-1}$ for $k \geq 3$. Ehrgott [8] considered a variant of MC-Min-STSP, where all objectives are encoded into a single objective by using some norm. He proved approximation ratios between 3/2 and 2 for this problem, where the ratio depends on the norm used.

Manthey and Ram [17] designed a $(2+\varepsilon)$ approximation algorithm for MC-Min-STSP and an approximation algorithm for MC-Min-γ-ATSP, which achieves a constant ratio but works only for $\gamma < 1/\sqrt{3} \approx 0.58$. They left open the existence of approximation algorithms for MC-Max-STSP, MC-Max-ATSP, and MC-Min-ATSP.

Bläser et al. [4] devised the first randomized approximation algorithms for MC-Max-STSP and MC-Max-ATSP. Their algorithms achieve ratios of $\frac{1}{k}+\varepsilon$ for k-C-Max-STSP and $\frac{1}{k+1}+\varepsilon$ for k-C-Max-ATSP. They argue that with their approach, only approximation ratios of $\frac{1}{k \pm O(1)}$ can be achieved, but they conjectured that ratios of $\Omega(1/\log k)$ are possible.

New Results. We devise approximation algorithms for MC-Max-STSP, MC-Max-ATSP, and MC-Min-ATSP. The approximation ratios achieved by our algorithms are independent of the number k of criteria, and they come close to the best approximation ratios known for Max-STSP, Max-ATSP, and Min-ATSP with only a single objective function. Our algorithms work for any number k of criteria.

First, we solve the conjecture of Bläser et al. [4] affirmatively. We even prove a stronger result: For MC-Max-STSP, we achieve a ratio of $2/3-\varepsilon$, while for MC-Max-ATSP, we achieve a ratio of $1/2-\varepsilon$ (Section 4). Already for $k=2$, this is an improvement from

$\frac{1}{2}-\varepsilon$ to $\frac{2}{3}-\varepsilon$ for 2-C-Max-STSP and from $\frac{1}{3}-\varepsilon$ to $\frac{1}{2}-\varepsilon$ for 2-C-Max-ATSP. The general idea of these algorithms is sketched in Section 2. After that, we introduce a decomposition technique in Section 3 that will lead to our algorithms. The running-time of our algorithms is polynomial in the input size for any fixed $\varepsilon > 0$ and any fixed number k of criteria.

Furthermore, we devise a *deterministic* approximation algorithm for 2-C-Max-STSP that achieves a ratio of $61/243 > 1/4$. As a side effect, this proves that for 2-C-Max-STSP, there always exists a single tour that already is a 1/3 approximate Pareto curve.

Finally, we devise the first approximation algorithm for MC-Min-ATSP (Section 6). In addition, our algorithm improves on the algorithm for MC-Min-γ-ATSP by Manthey and Ram [17] for $\gamma > 0.55$, and it is the first approximation algorithm for MC-Min-γ-ATSP for $\gamma \in [1/\sqrt{3}, 1)$. The approximation ratio of our algorithm is $\log n + \varepsilon$ for MC-Min-ATSP, where n is the number of vertices. Furthermore, it is a $\frac{1}{1-\gamma} + \varepsilon$ approximation for MC-Min-γ-ATSP for $\gamma \in [\frac{1}{2}, 1)$. Our algorithm is randomized.

Due to lack of space, most proofs are omitted. For complete proofs, we refer to the full version of this paper [16].

2. Outline and Idea for MC-Max-TSP

For Max-ATSP, we can easily get a 1/2 approximation: We compute a maximum-weight cycle cover, and remove the lightest edge of each cycle. In this way, we obtain a collection of paths. Then we add edges to connect the paths, which yields a Hamiltonian cycle. For Max-STSP, this approach yields a ratio of 2/3.

Unfortunately, this does not generalize to multi-criteria Max-TSP, even though $(1-\varepsilon)$ approximate Pareto curves of cycle covers can be computed in polynomial time. The reason is that the term “lightest edge” is usually not well defined: An edge that has little weight with respect to one objective might have a huge weight with respect to another objective. Based on this observation, the basic idea behind our algorithms is the following case distinction: First, if every edge of a cycle cover is a *light-weight edge*, i.e., it contributes only little to the overall weight, then removing one edge does not decrease the total weight by too much. We can choose the edges for removal carefully to get an approximate tour.

Second, if there is one edge that is very heavy with respect to one objective (a *heavy-weight edge*), then we contract this edge. We repeat this process until either we have obtained a cycle cover that contains only light-weight edges or we have enough weight for one objective. In the former case, we can use decomposition. In the latter case, we proceed recursively on the remaining graph with $k-1$ objectives.

In Section 3, we deal with the first case. This includes the definition of when we call an edge a light-weight edge. In Section 4, we present our algorithms, which includes the recursion in case of a heavy-weight edge. The approximation ratios that we achieve come close, i.e., up to an arbitrarily small additive $\varepsilon > 0$, to the 1/2 and 2/3 mentioned above for mono-criterion Max-ATSP and Max-STSP.

3. Decompositions

From any collection P of paths, we obtain a Hamiltonian cycle just by connecting the endpoints of the paths appropriately. Assume that we are given a cycle cover C. If we can find a collection of paths $P \subseteq C$ (by removing one edge of every cycle of C) with

$w(P) \geq \alpha \cdot w(C)$ for some $\alpha \in (0,1]$, then this would yield an approximate solution for Max-TSP. We call these paths P an *α-decomposition of C* for some $\alpha \in (0,1]$ if $w(P) \geq \alpha w(C)$. Not every cycle cover possesses an α-decomposition for every α. Let $k \geq 1$ be the number of criteria. Bläser et al. defined $\alpha_k^d \in [0,1]$ to be the maximum number such that the following holds: every directed cycle cover C with edge weights $w = (w_1, \ldots, w_k)$ that satisfies $w(e) \leq \alpha_k^d \cdot w(C)$ for all $e \in C$ possesses an α_k^d-decomposition. The value $\alpha_k^u \in [0,1]$ is analogously defined for undirected cycle covers. We have $\alpha_1^d = \frac{1}{2}$ and $\alpha_1^u = \frac{2}{3}$. We also have $\alpha_k^u \geq \alpha_k^d$ and $\alpha_k^u \leq \alpha_{k-1}^u$ as well as $\alpha_k^d \leq \alpha_{k-1}^d$.

Bläser et al. [4] proved $\alpha_k^d \geq \frac{1}{k+1}$ and $\alpha_k^u \geq \frac{1}{k}$. Furthermore, they proved the existence of $\Omega(1/\log k)$-decompositions, i.e., $\alpha_k^d, \alpha_k^u \in \Omega(1/\log k)$, which led to their conjecture that $\Omega(1/\log k)$ approximation algorithms might exist. However, their approximation algorithms do not make use of the $\Omega(1/\log k)$ decompositions, and they only achieve ratios of $\frac{1}{k} - \varepsilon$ for k-C-Max-STSP and $\frac{1}{k+1} - \varepsilon$ for k-C-Max-ATSP. In fact, they indicate that approximation ratios of $\frac{1}{k+O(1)}$ are the best that can be proved using their approach. For completeness, we make their decomposition result more precise with the next theorem. In particular, we show that $\alpha_k^d, \alpha_k^u \in \Theta(1/\log k)$, which proves that better approximations require a different decomposition technique. The new decompositions will be introduced later on in this section.

Theorem 3.1. *For all $1 \leq k \in \mathbb{N}$, we have*

$$\frac{1}{0.78 \cdot \log_2 k + \frac{3}{2}} \approx \frac{1}{\frac{9}{8} \cdot \ln k + \frac{3}{2}} \leq \alpha_k^u \leq \frac{1}{\lfloor \log_3 k \rfloor + 1} \approx \frac{1}{0.63 \cdot \log_2 k + 1} \text{ and}$$
$$\frac{1}{1.39 \cdot \log_2 k + 4} \approx \frac{1}{2 \cdot \ln k + 4} \leq \alpha_k^d \leq \frac{1}{\lfloor \log_2 k \rfloor + 2}.$$

In order to obtain constant approximation ratios, independent of k, we have to generalize the concept of decompositions. Let C be a cycle cover, and let $w = (w_1, \ldots, w_k)$ be edge weights. We say that the pair (C, w) is γ-light for some $\gamma \geq 1$ if $w(e) \leq w(C)/\gamma$ for all $e \in C$. In the following, let $\eta_{k,\varepsilon} = \frac{\varepsilon^2}{2 \ln k}$.

Theorem 3.2. *Let ε be arbitrary with $0 < \varepsilon < 1/2$, and let $k \geq 2$ be arbitrary. Let C be a cycle cover, and let $w = (w_1, \ldots, w_k)$ be edge weights such that (C, w) is $1/\eta_{k,\varepsilon}$-light. If C is directed, then there exists a collection $P \subseteq C$ of paths with $w(P) \geq (\frac{1}{2} - \varepsilon) \cdot w(C)$. If C is undirected, then there exists a collection $P \subseteq C$ of paths with $w(P) \geq (\frac{2}{3} - \varepsilon) \cdot w(C)$.*

We know that decompositions exist due to Theorem 3.2. But, in order to use them in approximation algorithms, we have to find them efficiently. In the remainder of this section, we devise devise a simple randomized algorithm for this job. There is also a deterministic algorithm that we call 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 $(\frac{1}{2} - \varepsilon)$- or $(\frac{2}{3} - \varepsilon)$-decomposition $P \subseteq C$, provided that (C, w) is $1/\eta_{k,\varepsilon}$-light. Due to lack of space, we do not describe Decompose here.

The randomized algorithm exploits Theorem 3.2: Assume that we have a cycle cover C with edge weights w such that (C, w) is $1/\eta_{k,\varepsilon}$-light. We randomly select one edge of every cycle of C for removal and put all remaining edges into P. The probability that P is not a $(\frac{1}{2} - \varepsilon)$- or $(\frac{2}{3} - \varepsilon)$-decomposition (depending on whether C is directed or undirected) is bounded from above by $1/k \leq 1/2$. Thus, we obtain a decomposition with constant probability. We iterate this process until a feasible decomposition has been found. In this way, we get a Las Vegas algorithm with expected linear running-time.

4. Approximation Algorithms for MC-Max-TSP

In this section, MaxCC-Approx denotes the randomized FPTAS for cycle covers. More precisely, let G be a graph (directed or undirected), $w = (w_1, \ldots, w_k)$ be edge weights, $\varepsilon > 0$ and $p \in (0, 1]$. Then MaxCC-Approx$(G, w, k, \varepsilon, p)$ yields a $(1 - \varepsilon)$-approximate Pareto curve of cycle covers of G with weights w with a success probability of at least $1 - p$.

4.1. Multi-Criteria Max-ATSP

Our goal is now either to use decomposition or to reduce the k-criteria instance to a $(k-1)$-criteria instance. To this aim, we put the cart before the horse: Instead of computing Hamiltonian cycles, we assume that they are given. Then we show how to force an algorithm to find approximations to them. To obtain a $1/2 - \varepsilon$ approximate Pareto curve, we have to make sure that for every tour $\tilde{H}$, we have a tour H in our set with $w(H) \geq (\frac{1}{2} - \varepsilon) \cdot w(\tilde{H})$. Fix ε with $0 < \varepsilon < \frac{1}{2 \ln k}$, let $\tilde{H}$ be any tour, and let $\beta_i = \max\{w_i(e) \mid e \in \tilde{H}\}$ be the weight of the heaviest edge with respect to the ith objective. Let $\beta = \beta(\tilde{H}) = (\beta_1, \ldots, \beta_k)$. We will distinguish two cases.

In the first case, we assume that $\beta \leq (\eta_{k,\varepsilon} - \varepsilon^3) \cdot w(\tilde{H})$, i.e., $\tilde{H}$ does not contain any heavy-weight edges. We modify our edge weights w to w^β as follows:

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

This means that we set all edge weights exceeding β to 0. Since $\tilde{H}$ does not contain any edges whose weight has been set to 0, we have $w(\tilde{H}) = w^\beta(\tilde{H})$. Furthermore, for all subsets C of edges, we have $w^\beta(C) \leq w(C)$. The advantage of w^β is that, if we compute a $(1 - \varepsilon)$ approximate Pareto curve $\mathcal{C}^\beta$ of cycle covers with edge weights w^β, we obtain a cycle cover to which we can apply decomposition to obtain a collection P of paths. Then P yields a tour H that approximates $\tilde{H}$. This is stated in the following lemma.

Lemma 4.1. *Let $\varepsilon > 0$ be arbitrary. Let $\tilde{H}$ be a directed tour with $w(e) \leq (\eta_{k,\varepsilon} - \varepsilon^3) \cdot w(\tilde{H})$ for all $e \in \tilde{H}$. Let $\beta = \beta(\tilde{H})$, and let $\mathcal{C}^\beta$ be a $(1 - \varepsilon)$ approximate Pareto curve of cycle covers with respect to w^β.*

Then $\mathcal{C}^\beta$ contains a cycle cover C that yields a decomposition $P \subseteq C$ with $w(P) \geq (\frac{1}{2} - 2\varepsilon) \cdot w(\tilde{H})$.

In the second case, we assume that there exists an edge $e = (u, v) \in \tilde{H}$ and an $i \in [k]$ with $w_i(e) > (\eta_{k,\varepsilon} - \varepsilon^3) \cdot w(\tilde{H})$. We put this edge into a set K of edges that we want to have in our cycle cover no matter what. Then we contract the edge e by removing all outgoing edges of u and all incoming edges of v and identifying u and v. In this way, we obtain a slightly smaller tour $\tilde{H}' = \tilde{H} \setminus \{e\}$. Again, there might be an edge $e' \in \tilde{H}'$ and an $i' \in [k]$ with $w_{i'}(e') > (\eta_{k,\varepsilon} - \varepsilon^3) \cdot w_{i'}(\tilde{H}')$. (Since $w(\tilde{H}') \leq w(\tilde{H})$, edges that have not been heavy can now be heavy with respect to $\tilde{H}'$.) We put e' into K, contract e' and recurse. How long can this process go on? There are two cases that can bring it to an end: First, we might get a tour H' that does not have any more heavy-weight edges, i.e., $w(e) \leq (\eta_{k,\varepsilon} - \varepsilon^3) \cdot w(H')$ for all $e \in H'$. In this case, we can apply Lemma 4.1 with decomposition. Second, we might get an $i \in [k]$ with $w_i(K) \geq (\frac{1}{2} - \varepsilon) \cdot w_i(\tilde{H})$, where $\tilde{H}$ is our original tour. Then we have collected enough weight with respect to the ith objective, and we can continue with only $k - 1$ objectives. The next lemma bounds the number of edges in K from above.

$\mathcal{P}_{\text{TSP}} \leftarrow$ MaxATSP-Approx$(G, w, k, \varepsilon, p)$
input: directed complete graph $G = (V, E)$, $k \geq 1$, edge weights $w : E \to \mathbb{N}^k$, $\varepsilon > 0$
output: $(\frac{1}{2} - \varepsilon)$ approximate Pareto curve $\mathcal{P}_{\text{TSP}}$ for k-C-Max-ATSP with a success probability of at least $1 - p$
if $k = 1$ **then**
  compute a 2/3 approximation $\mathcal{P}_{\text{TSP}}$
else
  for all subsets $K \subseteq E$ with $|K| \leq f(k, \varepsilon/2)$ such that K is a path cover **do**
    contract all edges of K to obtain G_K
    for all bounds β of (G_K, w) **do**
      $\mathcal{C}_{K,\beta} \leftarrow$ MaxCC-Approx$\big(G_K, w^\beta, k, \frac{\varepsilon}{2}, \frac{p}{2n^{2k+2f(k,\varepsilon/2)}}\big)$
      for all $C \in \mathcal{C}_{K,\beta}$ with $w^\beta(e) \leq \eta_{k,\varepsilon/2} \cdot w^\beta(C)$ for all $e \in C$ **do**
        $P \leftarrow$ Decompose$(C, w^\beta, \varepsilon/2)$
        add edges to $K \cup P$ to obtain a tour H; add H to $\mathcal{P}_{\text{TSP}}$
    for all $i \leftarrow 1$ **to** k **do**
      remove the ith objective from w to obtain w'
      $\mathcal{P}^{K,i}_{\text{TSP}} \leftarrow$ MaxATSP-Approx$(G_K, w', k-1, \frac{\varepsilon}{2}, \frac{p}{2n^{2k+2f(k,\varepsilon/2)}})$
      for all $H' \in \mathcal{P}^{K,i}_{\text{TSP}}$ **do**
        $H \leftarrow K \cup H'$; add H to $\mathcal{P}_{\text{TSP}}$

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

Lemma 4.2. *After at most* $f(k, \varepsilon) = k \cdot \left\lceil \frac{\log(1/2+\varepsilon)}{\log(1-\eta_{k,\varepsilon}+\varepsilon^3)} \right\rceil$ *iterations, the procedure described above halts.*

To obtain an algorithm, we have to find β and K. So far, we have assumed that we already know the Hamiltonian cycles that we aim for. But there is only a polynomial number of possibilities for β and K: For all β, we can assume that for all i there is an edge with $w_i(e) = \beta_i$. Thus, for every i there are at most $O(n^2)$ choices for β_i, hence at most $O(n^{2k})$ in total. The cardinality of K is bounded in terms of $f(k, \varepsilon)$. For fixed k and ε, there is only a polynomial number of subsets of cardinality at most $f(k, \varepsilon)$. We can even restrict ourselves to the subsets K that are path covers: A path cover is a subset K of edges such that K does not contain cycles and both the indegree and outdegree of every vertex is at most one. We obtain MaxATSP-Approx (Alg. 1) and the following theorem.

Theorem 4.3. *For every* $k \geq 1$, $\varepsilon > 0$, MaxATSP-Approx *is a randomized* $\frac{1}{2} - \varepsilon$ *approximation for k-criteria Max-ATSP whose running-time for a success probability of at least* $1 - p$ *is polynomial in the input size and* $\log(1/p)$.

Proof. We have to estimate three things: approximation ratio, running-time, and success probability. The proof is by induction on k. For $k = 1$, the theorem holds since there is a deterministic, polynomial-time 2/3 approximation for mono-criterion Max-ATSP. In the following, we assume that the theorem is correct for $k - 1$.

Let us focus on the approximation ratio, the other aspects are omitted for lack of space. For this purpose, we assume that all randomized computations are successful. Let $\tilde{H}$ be an arbitrary tour. For a subset $K \subseteq \tilde{H}$, let $\tilde{H}_K$ be $\tilde{H}$ with all edges in K being contracted. Then, by Lemma 4.2, there exists a (possibly empty) set $K \subseteq \tilde{H}$ of edges of cardinality at most $f(k, \varepsilon/2)$ with one of the two following properties:

(1) There exists an i with $w_i(K) \geq (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot w_i(\tilde{H})$.
(2) For all $e \in \tilde{H}_K$, we have $w(e) \leq (\eta_{k,\varepsilon/2} - (\frac{\varepsilon}{2})^3) \cdot w(\tilde{H}_K)$.

In the first case, there exists an $H' \in \mathcal{P}_{\text{TSP}}^{K,i}$ (see line 13) with $w_j(H') \geq (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot w_j(\tilde{H}_K)$ for all $j \in [k] \setminus \{i\}$. H' combined with K yields a tour H that satisfies $w(H) \geq (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot w(\tilde{H})$: First, we have $w_i(H) \geq w_i(K) \geq \left(\frac{1}{2} - \frac{\varepsilon}{2}\right) \cdot w_i(\tilde{H})$. Second, for $j \neq i$, we have $w_j(H) = w_j(K) + w_j(H') \geq w_j(K) + \left(\frac{1}{2} - \frac{\varepsilon}{2}\right) \cdot w_j(\tilde{H}_K) = \left(\frac{1}{2} - \frac{\varepsilon}{2}\right) \cdot w_j(\tilde{H})$.

In the second case, let $\beta_i = \max\{w_i(e) \mid e \in \tilde{H}_K\} \leq \left(\eta_{k,\varepsilon/2} - (\frac{\varepsilon}{2})^3\right) \cdot w_i(\tilde{H}_K)$. Then $\mathcal{C}_{K,\beta}$ contains a cycle cover C with $w(C) \geq (1 - \frac{\varepsilon}{2}) \cdot w(\tilde{H}_K)$ and $w^\beta(e) \leq \eta_{k,\varepsilon/2} \cdot w(\tilde{H}_K)$ (Lemma 4.1). Thus, C can be decomposed into a collection P of paths with $w(P) \geq (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot w(C)$ (Lemma 4.1). Together with K, this yields a tour H with

$$\begin{aligned} w(H) &\geq w(P) + w(K) \geq \left(\tfrac{1}{2} - \tfrac{\varepsilon}{2}\right) \cdot w(C) + w(K) \geq \left(\tfrac{1}{2} - \tfrac{\varepsilon}{2}\right) \cdot \left(1 - \tfrac{\varepsilon}{2}\right) \cdot w(\tilde{H}_K) + w(K) \\ &= \left(\tfrac{1}{2} - \tfrac{3\varepsilon}{4} + \tfrac{\varepsilon^2}{4}\right) \cdot w(\tilde{H}_K) + w(K) \geq \left(\tfrac{1}{2} - \varepsilon\right) \cdot w(\tilde{H}). \end{aligned}$$

∎

4.2. Multi-Criteria Max-STSP

The approximation for MC-Max-ATSP works of course also for MC-Max-STSP. Our goal, however, is a ratio of $(\frac{2}{3} - \varepsilon)$. As a first attempt, one might just replace the $(\frac{1}{2} - \varepsilon)$-decompositions by $(\frac{2}{3} - \varepsilon)$-decompositions. Unfortunately, this is not sufficient since contracting the heavy-weight edges in undirected graphs is not as easy as it is for directed graphs: First, both statements "remove all incoming" and "remove all outgoing" edges are not well-defined in an undirected graph. Second, if we just consider all edges of one vertex as the incoming edges and all edges of the other vertex as the outgoing edges, we obtain a directed graph, which allows only for a ratio of $\frac{1}{2} - \varepsilon$. To circumvent these problems, we do not contract edges $e = \{u, v\}$. Instead, we set the weight of all edges incident to u or v to 0. This allows us to add the edge e to any tour H' without decreasing the weight: We remove all edges incident to u or v from H', and then we add e. The result is a collection of paths. Then we add edges to connect these paths to a Hamiltonian cycle. The only edges that we have removed are edges incident to u or v, which have weight 0 anyway.

However, by setting the weight of edges adjacent to u or v to 0, we might destroy a lot of weight with respect to some objective. To solve this problem as well, we consider larger neighborhoods of the edges in K. In this way, we can add our heavy-weight edge (plus some more edges of its neighborhood) to the Hamiltonian cycle without losing too much weight from removing other edges. The function h in Alg. 2 depends only on k and ε and plays a similar role as f in Section 4.1. We omit the details and obtain MaxSTSP-Approx (Alg. 2) and the following theorem.

Theorem 4.4. *For every $k \geq 1$, $\varepsilon > 0$,* MaxSTSP-Approx *is a randomized $\frac{2}{3} - \varepsilon$ approximation for k-criteria Max-STSP whose running-time for a success probability of at least $1 - p$ is polynomial in the input size and* $\log(1/p)$.

$\mathcal{P}_{\text{TSP}} \leftarrow$ MAXSTSP-APPROX$(G, w, k, \varepsilon, p)$
input: undirected complete graph $G = (V, E)$, $k \geq 1$, edge weights $w : E \to \mathbb{N}^k$, $\varepsilon > 0$
output: $(\frac{2}{3} - \varepsilon)$ approximate Pareto curve $\mathcal{P}_{\text{TSP}}$ for k-C-Max-ATSP with a success probability of at least $1 - p$
if $k = 1$ **then**
    compute a 61/81 approximation $\mathcal{P}_{\text{TSP}}$
else
    for all subsets $K \subseteq E$ with $|K| \leq h(k, \varepsilon/3)$ such that K is a path cover **do**
        let L be the set of vertices incident to K
        obtain w^L from w by setting the weight of all edges incident to L to 0
        for all bounds β of (G, w^L) **do**
            $\mathcal{C}_{L,\beta} \leftarrow$ MAXCC-APPROX$\big(G, w^{L\beta}, k, \frac{\varepsilon}{3}, \frac{p}{2n^{2k+2h(k,\varepsilon/3)}}\big)$
            for all $C \in \mathcal{C}_{L,\beta}$ with $w^{L\beta}(e) \leq \eta_{k,\varepsilon/3} \cdot w^{L\beta}(C)$ for all $e \in C$ **do**
                $P \leftarrow$ DECOMPOSE$(C, w^{L\beta}, \varepsilon/3)$; remove edges of weight 0 from P
                add edges to $K \cup P$ to obtain a tour H; add H to $\mathcal{P}_{\text{TSP}}$
        for all $i \leftarrow 1$ **to** k **do**
            remove the ith objective from w^L to obtain w'^L
            $\mathcal{P}_{\text{TSP}}^{L,i} \leftarrow$ MAXATSP-APPROX$(G, w'^L, k-1, \frac{\varepsilon}{3}, \frac{p}{2n^{2k+2h(k,\varepsilon/3)}})$
            for all $H' \in \mathcal{P}_{\text{TSP}}^{L,i}$ **do**
                remove edges of weight 0 from H'
                add edges to $H' \cup K$ to obtain a tour H; add H to $\mathcal{P}_{\text{TSP}}$

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

5. Deterministic Approximations for 2-C-Max-STSP

The algorithms presented in the previous section are randomized due to the computation of approximate Pareto curves of cycles covers. So are most approximation algorithms for multi-criteria TSP. As a first step towards deterministic approximation algorithms for MC-Max-TSP, we present a deterministic $61/243 \approx 0.251$ approximation for 2-C-Max-STSP. The key insight for the results of this section is the following lemma, which yields tight bounds for the existence of approximate Pareto curves with only a single element (Theorem 5.2). For completeness, we note that single-element approximate Pareto curves exist for no other variant of multi-criteria TSP than 2-C-Max-STSP.

Lemma 5.1. *Let M be a matching, let H be a collection of paths or a Hamiltonian cycle, and let w be edge weights. Then there exists a subset $P \subseteq H$ such that $P \cup M$ is a collection of paths or a Hamiltonian cycle (we call P in this case an M-feasible set) and $w(P) \geq w(H)/3$.*

Theorem 5.2. *For every undirected complete graph G with edge weights w_1 and w_2, there exists a tour H such that $\{H\}$ is a $1/3$ approximate Pareto curve for 2-C-Max-STSP. This is tight: There exists a graph G with edge weights w_1 and w_2 such that, for all $\varepsilon > 0$, no single Hamiltonian tour of G is a $(1/3 + \varepsilon)$ approximate Pareto curve.*

Lemma 5.1 and Theorem 5.2 are constructive in the sense that, given a tour H_2 that maximizes w_2, the tour H can be computed in polynomial time. A matching M with $w_1(M) \geq w_1(H_1)/3$ can be computed in cubic time. However, since we cannot compute an optimal H_2 efficiently, the results cannot be exploited directly to get an algorithm. Instead, we use an approximation algorithm for finding a tour with as much weight with respect

$\mathcal{P}_{\text{TSP}} \leftarrow$ BiMaxSTSP-Approx(G, w_1, w_2)
input: undirected complete graph $G = (V, E)$, edge weights $w_1, w_2 : E \to \mathbb{N}^k$
output: a 61/243 approximate Pareto curve H
compute a maximum-weight matching M with respect to w_1
compute a 61/81 approximate tour H_2 with respect to w_2
$P \leftarrow H_2 \cap M$; $M' \leftarrow M$; $H_2 \leftarrow H_2 \setminus P$
while $H_2 \neq \emptyset$ **do**
    $e \leftarrow \text{argmax}\{w_2(e') \mid e' \in H_2\}$
    extend e to a path $e_1, \ldots, e_q \in H_2$ such that only e_1 and e_q are incident to edges $z_1, z_2 \in M'$ or the path cannot be extended anymore
    $P \leftarrow P \cup \{e_1, \ldots, e_q\}$; $H_2 \leftarrow H_2 \setminus \{e_1, \ldots, e_q\}$
    if z_1 or z_2 exists **then**
        let $f_1, f_2 \in H_2$ be the two edges extending the path if they exist
        $H_2 \leftarrow H_2 \setminus \{f_1, f_2\}$
        if both z_1 and z_2 exist **then** contract z_1 and z_2 to z; $M' \leftarrow (M' \setminus \{z_1, z_2\}) \cup \{z\}$
let H be a tour obtained from $P \cup M$

Algorithm 3: Approximation algorithm for 2-C-Max-STSP.

to w_2 as possible. Using the 61/81 approximation algorithm for Max-STSP [7], we obtain Alg. 3 and the following theorem.

Theorem 5.3. BiMaxSTSP-Approx *is a deterministic* 61/243 *approximation algorithm with running-time* $O(n^3)$ *for* 2*-C-Max-STSP.*

For metric 2-C-Max-STSP, i.e., the edge weights have to fulfil the triangle inequality, we obtain the an approximation ratio of $7/24 > 0.29$ if we replace the 61/81 approximation with the 7/8 approximation for metric Max-STSP by Kowalik and Mucha [15].

6. Approximation Algorithm for MC-Min-ATSP

Now we turn to MC-Min-ATSP and MC-Min-γ-ATSP, i.e., tours of *minimum* weight are sought in directed graphs. Alg. 4 is an adaptation of the algorithm of Frieze et al. [12] to multi-criteria ATSP. Therefore, we briefly describe their algorithm: We compute a cycle cover of minimum weight. If this cycle cover is already a Hamiltonian cycle, then we are done. Otherwise, we choose an arbitrary vertex from every cycle. Then we proceed recursively on the subset of vertices thus chosen to obtain a tour that contains all these vertices. The cycle cover plus this tour form an Eulerian graph. We traverse the Eulerian cycle and take shortcuts whenever we visit vertices more than once. The approximation ratio achieved by this algorithm is $\log_2 n$ for Min-ATSP [12] and $1/(1-\gamma)$ for Min-γ-ATSP [5].

MinATSP-Approx (Alg. 4) for MC-Min-ATSP proceeds as follows: MinCC-Approx computes an approximate Pareto curve of cycle covers. (MinCC-Approx$(G, w, k, \varepsilon, p)$ computes a $(1+\varepsilon)$ approximate Pareto curve of cycle covers of G with weights w with a success probability of at least $1-p$ in time polynomial in the input size, $1/\varepsilon$, and $\log(1/p)$.) Then we iterate by computing approximate Pareto curves of cycle covers on vertex sets V' for every cycle cover C in the previous set. The set V' contains exactly one vertex of every cycle of C. Unfortunately, it can happen that we construct a super-polynomial number of solutions in this way. To cope with this, we remove some intermediate solutions if there are other intermediate solutions whose weight is close by. We call this process *sparsification*.

$\mathcal{P}_{\text{TSP}} \leftarrow$ MinATSP-Approx(G, w, k, ε)
input: directed complete graph $G = (V, E)$ with $n = |V|$, $k \geq 1$, $w : E \to \mathbb{N}^k$, $\varepsilon > 0$
output: $(\log n + \varepsilon)$ approximate Pareto curve for k-C-Min-ATSP or $(\frac{1}{1-\gamma} + \varepsilon)$ approximate Pareto curve for k-C-Min-γ-ATSP with a probability of at least 1/2

1: $\varepsilon' \leftarrow \varepsilon^2 / \log^3 n$; $\mathcal{F} \leftarrow \emptyset$; $j \leftarrow 1$
2: $\mathcal{C} \leftarrow$ MinCC-Approx$\big(G, w, k, \varepsilon', \frac{1}{2Q \log n}\big)$
3: $\mathcal{P}_0 \leftarrow \{(C, w(C), V, \perp) \mid C \in \mathcal{C}\}$
4: **while** $\mathcal{P}_{j-1} \neq \emptyset$ **do**
5:     $\mathcal{P}_j \leftarrow \emptyset$
6:     **for all** $\pi = (C', w', V', \pi') \in \mathcal{P}_{j-1}$ **do**
7:         **if** (V', C') is connected **then** $\mathcal{F} \leftarrow \mathcal{F} \cup \{(C', w', V', \pi')\}$
8:         **else**
9:             select one vertex of every component of (V', C') to obtain $\tilde{V}$
10:             $\mathcal{C} \leftarrow$ MinCC-Approx$\big(G, w, k, \varepsilon', \frac{1}{2Q \log n}\big)$
11:             $\mathcal{P}_j \leftarrow \mathcal{P}_j \cup \{(\tilde{C}, \tilde{w}, \tilde{V}, \pi) \mid \tilde{C} \in \mathcal{C}, \tilde{w} = w' + \gamma^j \cdot w(\tilde{C})\}$
12:     **while** there are $\pi', \pi'' \in \mathcal{P}_j$ with equal ε'-signature **do** remove one of them
13:     $j \leftarrow j + 1$
14: $\mathcal{P}_{\text{TSP}} \leftarrow \emptyset$
15: **for all** $(C', w', V', \pi') \in \mathcal{F}$ **do**
16:     $H \leftarrow C'$
17:     **while** $\pi' = (C'', w'', V'', \pi'') \neq \perp$ **do**
18:         construct tour H' on V'' from $H \cup C''$ by taking shortcuts such that $H \cap H' = \emptyset$
19:         $\pi' \leftarrow \pi''$; $H \leftarrow H'$
20:     $\mathcal{P}_{\text{TSP}} \leftarrow \mathcal{P}_{\text{TSP}} \cup \{H\}$

Algorithm 4: Approximation algorithm for MC-Min-ATSP and MC-Min-γ-ATSP.

It is based on the following observation: Let $\varepsilon > 0$, and consider H of weight $w(H) \in \mathbb{N}^k$. For every $i \in \{1, \ldots, k\}$, there is a unique $\ell_i \in \mathbb{N}$ such that $w_i(H) \in [(1+\varepsilon)^{\ell_i}, (1+\varepsilon)^{\ell_i+1})$. We call the vector $\ell = (\ell_1, \ldots, \ell_k)$ the *ε-signature* of H and of $w(H)$. Since $w(H) \leq 2^{p(N)}$, there are at most q^k different ε-signatures for some polynomial k, which is polynomial for fixed k. To get an approximate Pareto curve, we can restrict ourselves to have at most one solution with any specific ε-signature.

In the loop in lines 4 to 13, MinATSP-Approx computes iteratively Pareto curves of cycle covers. The set $\mathcal{P}_j$ contains *configurations* $\pi = (C', w', V', \pi')$, where C' is a cycle cover on V', π' is the predecessor configuration, and w' is the weight of C' plus the weight of its predecessor cycle covers, each weighted with an appropriate power of γ. (We define the ε'-signature of $\pi = (C', w', V', \pi')$ to be the ε'-signature of w'.) These weights are needed for the analysis of the approximation ratio. If, in the course of this computation, we obtain Hamiltonian cycles, these are put into $\mathcal{F}$ (line 7). In line 12, the sparsification takes place. Finally, in lines 14 to 20, Hamiltonian cycles are constructed from the cycle covers computed. In the algorithm, $Q = Q(N, 1/\varepsilon')$ is a two-variable polynomial that bounds the number of different ε'-signatures of solutions for instances of size at most N.

MinATSP-Approx is the first approximation algorithm for MC-Min-ATSP and for MC-Min-γ-ATSP for $\gamma \geq \sqrt{1/3} \approx 0.58$. For $\gamma > 0.55$, it improves over the previously known algorithm [17], which achieves a ratio of $\frac{1}{2} + \frac{\gamma^3}{1-3\gamma^2}$ and works only for $\gamma < \sqrt{1/3} \approx 0.58$.

Theorem 6.1. *For every $\varepsilon > 0$, Alg. 4 is a randomized $(\log n + \varepsilon)$ approximation for MC-Min-ATSP and a randomized $(\frac{1}{1-\gamma} + \varepsilon)$ approximation for MC-Min-γ-ATSP for $\gamma \in [\frac{1}{2}, 1)$. Its running-time is polynomial in the input size and $1/\varepsilon$.*

7. Open Problems

Most approximation algorithms for multi-criteria TSP use randomness for computing approximate Pareto curves of cycle covers. This raises the question if there are algorithms for multi-criteria TSP that are faster, deterministic, and achieve better approximation ratios.

References

[1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows*. Prentice-Hall, 1993.

[2] 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.

[3] Eric Angel, Evripidis Bampis, Laurent Gourvès, and Jérọme Monnot. (Non-)approximability for the multi-criteria TSP(1,2). In *Proc. 15th Int. Symp. on Fundamentals of Computation Theory (FCT)*, vol. 3623 of *LNCS*, pp. 329–340. Springer, 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 *LNCS*, pp. 185–197. Springer, 2008.

[5] Markus Bläser, Bodo Manthey, and Jiří Sgall. An improved approximation algorithm for the asymmetric TSP with strengthened triangle inequality. *J. Discrete Algorithms*, 4(4):623–632, 2006.

[6] L. Sunil Chandran and L. Shankar Ram. On the relationship between ATSP and the cycle cover problem. *Theoret. Comput. Sci.*, 370(1-3):218–228, 2007.

[7] Zhi-Zhong Chen, Yuusuke Okamoto, and Lusheng Wang. Improved deterministic approximation algorithms for Max TSP. *Inform. Process. Lett.*, 95(2):333–342, 2005.

[8] Matthias Ehrgott. Approximation algorithms for combinatorial multicriteria optimization problems. *Int. Trans. Oper. Res.*, 7(1):5–31, 2000.

[9] Matthias Ehrgott. *Multicriteria Optimization*. Springer, 2005.

[10] Matthias Ehrgott and Xavier Gandibleux. A survey and annotated bibliography of multiobjective combinatorial optimization. *OR Spectrum*, 22(4):425–460, 2000.

[11] 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 *LNCS*, pp. 104–118. Springer, 2007.

[12] Alan M. Frieze, Giulia Galbiati, and Francesco Maffioli. On the worst-case performance of some algorithms for the traveling salesman problem. *Networks*, 12(1):23–39, 1982.

[13] Paul C. Gilmore, Eugene L. Lawler, and David B. Shmoys. Well-solved special cases. In Eugene L. Lawler et al., editors, *The Traveling Salesman Problem*, pp. 87–143. John Wiley & Sons, 1985.

[14] 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.

[15] Łukasz Kowalik and Marcin Mucha. Deterministic 7/8-approximation for the metric maximum TSP. In *Proc. 11th Int. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, vol. 5171 of *LNCS*, pp. 132–145. Springer, 2008.

[16] Bodo Manthey. On approximating multi-criteria TSP. CoRR 0711.2157 [cs.DS], arXiv, 2008.

[17] Bodo Manthey and L. Shankar Ram. Approximation algorithms for multi-criteria traveling salesman problems. *Algorithmica*, to appear.

[18] 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 Computer Society, 2000.

This work is licensed under the Creative Commons Attribution-NoDerivs License. To view a copy of this license, visit `http://creativecommons.org/licenses/by-nd/3.0/`.