

A $(3/2 + \varepsilon)$ -Approximation for Multiple TSP with a Variable Number of Depots

Max Deppert  

Institute for Algorithms and Complexity, Hamburg University of Technology, Germany

Matthias Kaul  

Institute for Algorithms and Complexity, Hamburg University of Technology, Germany

Matthias Mnich  

Institute for Algorithms and Complexity, Hamburg University of Technology, Germany

Abstract

One of the most studied extensions of the famous Traveling Salesperson Problem (TSP) is the MULTIPLE TSP: a set of $m \geq 1$ salespersons collectively traverses a set of n cities by m non-trivial tours, to minimize the total length of their tours. This problem can also be considered to be a variant of UNCAPACITATED VEHICLE ROUTING, where the objective is to minimize the sum of all tour lengths. When all m tours start from and end at a *single* common *depot* v_0 , then the metric MULTIPLE TSP can be approximated equally well as the standard metric TSP, as shown by Frieze (1983).

The metric MULTIPLE TSP becomes significantly harder to approximate when there is a *set* D of $d \geq 1$ depots that form the starting and end points of the m tours. For this case, only a $(2 - 1/d)$ -approximation in polynomial time is known, as well as a $3/2$ -approximation for *constant* d which requires a prohibitive run time of $n^{\Theta(d)}$ (Xu and Rodrigues, *INFORMS J. Comput.*, 2015). A recent work of Traub, Vygen and Zenklusen (STOC 2020) gives another approximation algorithm for metric MULTIPLE TSP with run time $n^{\Theta(d)}$, which reduces the problem to approximating metric TSP.

In this paper we overcome the $n^{\Theta(d)}$ time barrier: we give the first efficient approximation algorithm for MULTIPLE TSP with a *variable* number d of depots that yields a better-than-2 approximation. Our algorithm runs in time $(1/\varepsilon)^{\mathcal{O}(d \log d)} \cdot n^{\mathcal{O}(1)}$, and produces a $(3/2 + \varepsilon)$ -approximation with constant probability. For the graphic case, we obtain a deterministic $3/2$ -approximation in time $2^d \cdot n^{\mathcal{O}(1)}$.

2012 ACM Subject Classification Theory of computation \rightarrow Approximation algorithms analysis

Keywords and phrases Traveling salesperson problem, rural postperson problem, multiple TSP, vehicle routing

Digital Object Identifier 10.4230/LIPIcs.ESA.2023.39

Funding Partially supported by DFG project MN 59/4-1.

Acknowledgements The third author thanks László Végh for inspiring discussions on multi-depot TSP and feedback on an earlier version.

1 Introduction

The TRAVELING SALESPERSON PROBLEM (TSP) is one of the best-studied problems in combinatorial optimization: given a complete graph G on n nodes together with edge weights $w : E(G) \rightarrow \mathbb{R}_{\geq 0}$, we seek a tour that starts at some node $v_0 \in V(G)$, then visits all other nodes of G exactly once, and returns to the origin v_0 in such a way that the overall tour weight is minimized, which is the sum of the weights of the edges traversed by the tour. TSP is one of Karp's 21 NP-complete problems [15], which motivates the design of efficient, polynomial-time approximation algorithms for it. Recall that an α -approximation for a minimization problem returns, for any instance I , in polynomial time a solution of value at



© Max Deppert, Matthias Kaul, and Matthias Mnich;
licensed under Creative Commons License CC-BY 4.0
31st Annual European Symposium on Algorithms (ESA 2023).

Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 39;
pp. 39:1–39:15



Leibniz International Proceedings in Informatics
LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany


most $\alpha \cdot \text{OPT}(I)$, where $\text{OPT}(I)$ denotes the value of an optimal solution for I . Of special importance in this regard is METRIC TSP, when the edge weight function w obeys the triangle inequality. For METRIC TSP, the tree doubling heuristic yields a 2-approximation, which was improved to a $3/2$ -approximation by Christofides [7] and Serdyukov [24] in the 1970s. This approximation factor stood unchallenged for many decades until its recent improvement to a $(3/2 - 10^{-36})$ -approximation by Karlin et al. [14].

Due to its ubiquity, a large variety of extensions of the TSP have been studied. Among the most prominent ones is the MULTIPLE TSP, where a set of $m \geq 1$ salespersons (all starting from some common node v_0 called a *depot*) jointly traverse the entire set of n nodes, in order to minimize the overall tour length. That is, the goal is to find a collection of m pairwise edge-disjoint cycles C_1, \dots, C_m (all intersecting in some node v_0) in G whose union covers all nodes of the graph and such that the sum of the weights of the cycles is minimized. This character of having to solve both a partitioning and a sequencing problem simultaneously gives rise to considerable added complexity, akin to that encountered in vehicle routing problems. Indeed, one could interpret this problem as a variant of the UNCAPACITATED VEHICLE ROUTING PROBLEM; we, however, will adhere to the TSP-style naming convention, since this is more prevalent in the literature. Let us just mention here that for metric edge weights, MULTIPLE TSP has the same approximation guarantee as the standard (single-person) metric TSP; in particular, METRIC MULTIPLE TSP admits a $(3/2 - 10^{-36})$ -approximation in polynomial time by the results of Karlin et al. [14]. Frieze [10] analysed the case of METRIC MULTIPLE TSP when each of the tours has to contain at least one edge and intersect a common depot v_0 ; he provided a $3/2$ -approximation for this setting in polynomial time. The MULTIPLE TSP is studied in more than 1,300 publications; an extensive survey is provided by Bektaş [1].

In this paper we study an extension of the MULTIPLE TSP, where a *set* $D \subseteq V(G)$ of nodes is distinguished as *depots*. Formally, the MULTI-DEPOT MULTIPLE TSP (MDMTSP) takes as input a complete graph G on n nodes together with edge weights $w : E(G) \rightarrow \mathbb{R}_{\geq 0}$, as well as a set $D \subseteq V(G)$ of $d = |D|$ depots and an integer $m \geq 1$ denoting the number of salespersons available. Now again we are seeking a set of m pairwise edge-disjoint cycles C_1, \dots, C_m in G whose union covers all nodes of the graph and such that the sum of the weights of the cycles is minimized, but in addition each cycle must contain some depot from D . Such set of cycles is an optimal solution for the MDMTSP instance, and we denote the value of some optimal solution by $\text{OPT}(G, D, w)$ (or simply OPT if the instance is clear from the context). The MDMTSP is motivated by several applications of high practical impact, like motion planning of a set of unmanned aerial vehicles [17, 21, 31] and the routing of service technicians where the technicians are leaving from multiple depots [20].

The theoretical aspects of MDMTSP have been studied in many research papers [1, 2, 3, 6, 13, 16, 25, 26, 28, 29, 30]. At this point, let us issue a word of caution. There are quite a few other varieties of (MDM)TSP considered in the literature, all subtly different from each other. For a compact overview of possible variations, there is the review paper of Bektaş [1]. For the scope of this paper, we consider the metric MDMTSP where the edge weights form a metric. This allows us to assume that $m = d$ throughout. This assumption is made for the following two reasons: on the one hand, the case $m > d$ is negligible as the objective function (the total weight of all tours) is invariant for multiple tours starting from a single depot (if weights satisfy the triangle inequality, it is easy to show that there is always an optimal solution in which at most one route will start and end at each depot). On the other hand, in the case $m \leq d$ we can try each selection of $d' = m$ depots by paying a multiplicative factor of $\binom{d}{m}$ in the run time only. Thus, any instance of metric MDMTSP is specified by a triple (G, D, w) , where G is a complete graph on n nodes, $D \subseteq V(G)$ is the set of depots, and $w : E(G) \rightarrow \mathbb{R}_{\geq 0}$ is a metric.

The polynomial-time approximability of metric MDMTSP is not fully understood. That there is a set D of depots (and not just a single depot v_0), each one of which must be visited by one of the tours, makes the approximability of the problem much harder compared to metric MULTIPLE TSP (i.e., the version without depots). The added complexity arises from the fact that we not only have to give a good order in which to visit nodes, as in the TSP, but we also have to partition the nodes appropriately. In particular, the Christofides-Serdyukov algorithm [7, 24] no longer yields a $3/2$ -approximation in this setting. The original analysis of Christofides' and Serdyukov's algorithms relies on all odd-degree nodes of some spanning structure F lying on the same tour, so a parity-correcting edge set J can be computed that weighs at most $\frac{1}{2}$ OPT. This fact is not available in the multi-depot setting, so in polynomial time we can only guarantee a 2-approximation by using the spanner F for J also. However, this only achieves a tight approximation ratio of $2 - \frac{1}{d}$ for the multi-depot setting, as shown by Xu et al. [30] (see Figure 1 for a version of their lower-bound example), because the matching can have weight $\frac{d-1}{d}$ OPT.



■ **Figure 1** An instance on which Algorithm 1 achieves approximation ratio arbitrarily close to 2. Square nodes are depots and all edges have unit weight. Dashed edges indicate a tour of length d , dotted edges a minimum CSF with weight $d - 1$ that requires a join of weight $d - 1$ to complete.


To avoid this issue, the constrained spanning forest needs to be rearranged such that there is again a matching of weight $\frac{1}{2}$ OPT, as in the work of Xu and Rodrigues [28] – this rearrangement though requires $n^{\Theta(d)}$ time.

Similarly, the algorithmic approaches to metric TSP based on solving a linear program (LP) are also unlikely to give α -approximation algorithms with $\alpha < 2$ for metric MDMTSP. To this end, consider the following multi-depot version of the subtour-elimination LP, MDMTSP-LP:

$$\begin{aligned}
 & \text{minimize} && \sum_{e \in E(G)} w_e x_e \\
 & \text{subject to} && \sum_{e \in \delta(v)} x_e = 2, && \forall v \in V(G) \setminus D \\
 & && \sum_{e \in \delta(U)} x_e \geq 2, && \forall U \subseteq V(G) \setminus D \\
 & && x_e \in [0, 2], && \forall e \in E(G)
 \end{aligned} \tag{MDMTSP-LP}$$

In Figure 2 we give a construction to show that MDMTSP-LP has integrality gap 2.

If one gives up on the polynomial run time of the approximation algorithm, then smaller approximation factors are possible. Xu and Rodrigues [28] show how to obtain a $3/2$ -approximation, but their algorithm requires time $n^{\Theta(d)}$, which is polynomial only if the number $d = |D|$ of depots is constant. Another $3/2$ -approximation for MDMTSP with run time $n^{\Theta(d)}$ follows from the recent work of Traub, Vygen and Zenklusen [26]. They in fact show the much stronger result that any λ -approximation algorithm for metric TSP also gives a $(\lambda + \varepsilon)$ -approximation algorithm for metric MDMTSP with an additional run time factor of $n^{\mathcal{O}(d/\varepsilon)}$. In summary, the state-of-the-art for metric MDMTSP is that there is no α -approximation known for MDMTSP for any absolute constant $\alpha < 2$ which runs in time $n^{o(d)}$.



■ **Figure 2** An instance on which the multi-depot subtour-elimination LP has integrality gap arbitrarily close to 2. The square nodes are the depots and all edges have unit weight. The dashed edges are assigned $x_e = 0.5$ by the LP, the other edges $x_e = 1$. The LP then has optimum value at most $\ell + 6$, whereas the MDMTSP has optimum value $2(\ell + 4)$.

1.1 Our Results

Our main result is a novel approximation algorithm for metric MULTIPLE TSP on d depots with a significantly improved run time. That is, we provide the first algorithm for metric MDMTSP which breaks the $n^{\Theta(d)}$ time barrier to obtain an approximation ratio strictly better than 2. Given an instance (G, D, w) of metric MDMTSP, we say a collection C_1, \dots, C_d of cycles is a *tour* if they jointly cover all nodes of G and each cycle contains exactly one depot from D .

► **Theorem 1.** *There is an algorithm that, given any $\varepsilon > 0$, in time $(1/\varepsilon)^{\mathcal{O}(d \log d)} \cdot n^{\mathcal{O}(1)}$ computes a tour T for any set of n cities with metric distances and d depots. The algorithm is randomized, and with constant probability the length of the tour T is at most $(3/2 + \varepsilon) \cdot \text{OPT}$.*

Thus, our result significantly improves on the previously best run time $n^{\Theta(d)}$ by Xu and Rodrigues [28] at the cost of some small additive ε in the approximation factor.

To break through the barrier of 2 on the approximation ratio, we need to rework the initial spanner F to be “correctly aligned” with the optimal solution so that each subtour contains an even number of odd-degree nodes, as initially proposed by Xu and Rodrigues [28]. We show that an approximate reworking can be done in time $f(d, \varepsilon) \cdot n^{\mathcal{O}(1)}$ for some suitable function f , resulting in a $(3/2 + \varepsilon)$ -approximation. To this end, firstly, we give a reduction of metric MDMTSP to a related routing problem which is known as the RURAL POSTPERSON PROBLEM (RPP). In the RPP, we are given an edge-weighted graph G and a set R of required edges, and are asked to compute a minimum-weight edge set F such that $R \cup F$ is connected and Eulerian. Our reduction reveals an approximation algorithm with run time $\mathcal{O}(n^3 + t)$ to compute solutions no worse than $\frac{3}{2} \text{OPT} + \frac{1}{2}w(T)$, where $w(T)$ is the weight of a single-person TSP tour T through the depots and t denotes the time to compute T . Then we use a randomized algorithm of Gutin et al. [11] for the RPP, and an approximate weight reduction scheme of van Bevern et al. [27], to construct a $(1 + \varepsilon)$ -approximation algorithm for a variant of RPP with depots.

We are then in a position to speed up the reworking of the initial spanner due to two key insights. Firstly, we allow for some misalignment to remain, as long as it is only due to the presence of some *light* edges, limiting the number of edges we have to consider for removal from F . Secondly, we employ the constructed approximation algorithm for the RPP to complete our now disconnected spanner to a tour. Doing this provides a large speedup over the algorithm of Xu and Rodrigues, who first need to guess a set of edges to reconnect their spanner, and then employ a matching algorithm to obtain a tour. Using the RPP allows us to do both of these steps simultaneously, and considerably faster.

An important special case of metric TSP is when the metric w is induced by the shortest paths in a graph. This version is also known as GRAPHIC TSP, and has been studied extensively from the perspective of approximation algorithms [18, 23]. For MDMTSP on graphic metrics, we obtain a *deterministic algorithm* with slightly better approximation factor and a reduced run time.

► **Theorem 2.** *There is an algorithm that, given any graph G on n nodes and set $D \subseteq V(G)$ of d depots, in time $2^d \cdot n^{\mathcal{O}(1)}$ computes a tour T of length at most $\frac{3}{2} \cdot \text{OPT}$.*

1.2 Related work

In the case of a single salesperson, i.e. $m = 1$, Bérczi et al. [4] gave a polynomial-time $3/2$ -approximation for the *many-visits* version of metric MDMTSP, that is, when each node v is equipped with a request $r(v)$ (encoded in binary) of how many times it should be visited.

In a different work, Bérczi et al. [5] have shown constant-factor approximation algorithms with ratio at most 4 for variants of metric MDMTSP where each tour has to visit exactly one depot (and thus, $d \leq m$).

For the RPP, which asks for a minimum-weight tour traversing all edges of a given subset R of edges of a graph, there is a polynomial-time approximation algorithm (cf. Frederickson [8] or Jansen [12]) similar to the approach of Christofides-Serdyukov for metric TSP. The weight of a solution can be bounded by $(3\text{OPT} + w(R))/2$, where $w(R)$ is the weight of R .

Oberlin et al. [19] studied heuristic approaches for MDMTSP where $d = m$ and each salesperson is located at its own depot.

For the objective of minimizing the longest tour length of any salesperson (rather than the sum of all the tour lengths), Frederickson et al. [9], among other routing problems, considered the case of a single depot ($d = 1$), and presented a $(\rho + 1 - 1/m)$ -approximation algorithm where ρ is the approximation ratio of an algorithm for the single-salesperson TSP.

2 Preliminaries

Let \mathcal{U} be a finite universe. For a function $w : \mathcal{U} \rightarrow \mathbb{R}$ and a multiset $U \subseteq \mathcal{U}$ we write $w(U)$ to mean $\sum_{u \in U} w(u)$, where the sum has an additional summand for each copy of an element in U , i.e. it considers multiplicities. The disjoint union $\dot{\bigcup}_i A_i$ of some sets $\{A_i\}_i$ is considered to be the multiset of all items in the collection. For brevity we often write $2A$ to mean $A \dot{\cup} A$.

Throughout this paper, we consider the multiple-depot version of the metric MULTIPLE TSP, or metric MDMTSP for short. We will generally represent the metric by an edge-weighted graph whose shortest-path metric we assume to be the metric in use. Notice that this makes no difference in our setting since we are allowed to traverse edges multiple times; only if this is forbidden does the non-metric case become relevant.

► **Definition 3.** *An instance (G, D, w) of metric MDMTSP consists of a complete graph G , a set $D \subseteq V(G)$ of d depots, and a metric $w : E(G) \rightarrow \mathbb{R}_+$ on $V(G)$. A multiset of edges $T \subseteq E(G)$ is called a tour of (G, D, w) if*

(P1) *the multigraph $(V(G), T)$ has even degree at every node in $V(G)$,*

(P2) *and each connected component of $(V(G), T)$ contains at least one node from D .*

We denote by $\text{OPT}(G, D, w)$ the minimum weight of any tour of (G, D, w) . If the instance is clear from the context, we may only say OPT .

Edge sets are generally allowed to be multisets, and graphs can have parallel edges.

Imitating the general framework of Christofides-Serdyukov [7, 24], we first compute an edge set F , called a constrained spanning forest (CSF), that ensures the connectivity property (P2). We then compute an additional set of edges J such that $F \dot{\cup} J$ has property (P1).

► **Definition 4.** *Let G be a graph and let $D \subseteq V(G)$. A constrained spanning forest in (G, D) is a set $F \subseteq E(G)$ of edges such that the graph $(V(G), F)$ is acyclic and every connected component of $(V(G), F)$ contains at least one node from D .*

We will make use of the following result.

► **Theorem 5** (Rathinam et al. [21]). *Given any graph G on n nodes and m edges with weights $w : E(G) \rightarrow \mathbb{R}$ and a set $D \subseteq V(G)$, a minimum-weight CSF of (G, D, w) can be computed in time $\mathcal{O}((n + m) \log n)$.*

Proof. The computation of a minimum-weight CSF for (G, D, w) can be reduced to computing a minimum-weight spanning tree in a graph G' with edge weights w' . The graph G' is obtained from G by adding a single root node r connected to every depot by an edge of some weight less than the weight all other edges in G . Kruskal's algorithm is guaranteed to choose these edges for the minimum-weight spanning tree in (G', w') , and after removing r we are left with a minimum-weight CSF for (G, D, w) . ◀

Traditionally, property (P1) is obtained by computing a minimum-weight matching on the odd-degree nodes of the CSF, as in Algorithm 1. However, this algorithm only achieves a

■ **Algorithm 1** Algorithm MULTI-DEPOT CHRISTOFIDES-SERDYUKOV.

Input: A metric MDMTSP instance (G, D, w) .

Output: A tour T with $w(T) \leq 2 \text{OPT}$.

- 1 Compute a minimum-weight CSF F for (G, D, w) ;
- 2 Let U be the set of nodes with odd degree in F ;
- 3 Compute a minimum-weight perfect matching M in $G[U]$;

Result: $T := F \dot{\cup} M$

tight approximation ratio of $2 - \frac{1}{d}$ for the multi-depot setting, as shown by Xu et al. [30] (see Figure 1 for a simplified version of their lower bound example), because the matching can have weight $\frac{d-1}{d} \text{OPT}$. To avoid needing such an expensive matching, the constrained spanning forest needs to be rearranged such that there is again a matching of weight $\frac{1}{2} \text{OPT}$, as in the work of Xu and Rodrigues [28].

3 Reducing Multi-Depot Multiple TSP to Rural Postperson Problem

In this section we show a reduction from the metric MDMTSP to the RPP. Recall that in the RPP there is a required set R of edges that a tour should traverse, rather than a set of nodes.

► **Definition 6** (Rural Postperson Problem). *An instance (G, R, w) of RPP consists of a graph G , a set $R \subseteq E(G)$ of required edges, and a metric weight function $w : E(G) \rightarrow \mathbb{R}_{\geq 0}$. A solution is multiset $J \subseteq E(G)$ for which $(V, R \dot{\cup} J)$ is Eulerian, and which has only one non-singleton connected component.¹ The weight of a solution J is $w(J) = \sum_{e \in J} w(e)$. The goal of RPP is to compute an optimal solution, which is a solution of minimum weight $\text{OPT}(G, R, w)$.*

There is a polynomial-time approximation algorithm for RPP [8, 12] which computes a solution $J \subseteq E(G)$ such that $w(R \dot{\cup} J) \leq \frac{3}{2} w(R \dot{\cup} J^*)$, i. e. $w(J) \leq \frac{3}{2} \text{OPT} + \frac{1}{2} w(R)$, where J^* is some optimal solution. Due to the first inequality and the unavoidable weight of R , the algorithm is known as a $3/2$ -approximation for RPP. This situation is very similar to the current approximation status of metric MDMTSP, where we can obtain a $3/2$ -approximation if we allow for some additional additive term. This observation motivates the following reduction from metric MDMTSP to RPP.

¹ This means that nodes not incident to any edge from R do not need to be visited by the computed tour. For metric cost functions, however, one can always reduce to the case where R spans G .

► **Observation 7.** *For each instance (G, D, w) of MDMTSP there is an instance (G', R, w') of RPP such that any solution to the RPP instance can be transformed in polynomial time into a solution to the MDMTSP instance of the same weight.*

Proof. First, compute any TSP tour S on the depots in G , that is, on $G[D]$. Then, for each node $v \in V \setminus D$, introduce a second node v' , as well as an edge $e_v = \{v, v'\}$, and set its weight to $w'(e_v) = 0$. For each edge $e \in E(G)$ set $w'(e) = w(e)$, and set R to be the union of S and two copies of each e_v . The any solution to the constructed instance of RPP corresponds to an MDMTSP tour for (G, D, w) of the same weight. ◀

Notice that this reduction, together with the $3/2$ -approximation for RPP, allows us to compute a solution to MDMTSP of weight at most $\frac{3}{2} \text{OPT} + \frac{1}{2}w(S)$. In particular, if all depots are pairwise close to each other this is already a better-than-2 approximation.

In Section 5 we will in some sense show a stronger result that there is also a (Turing) reduction from MDMTSP to the special case of RPP where (V, R) has few connected components, which has been shown by Gutin et al. to be tractable [11]:

► **Proposition 8** (Gutin et al. [11]). *There is a randomized algorithm for RPP that for any instance (G, R, w) , where $(V(G), R)$ has k connected components and w takes only integer values, in time $2^{\mathcal{O}(k)}(n + \text{OPT}(G, R, w))^{\mathcal{O}(1)}$ produces a solution. With constant probability, the computed solution is optimal.*

However, as our reduction is only $(1+\varepsilon)$ -approximate with respect to the solution qualities, and needs time exponential in d and ε , we need to remove the polynomial dependence on $\text{OPT}(G, R, w)$ in Proposition 8. To this end, we will adapt an approximate weight reduction scheme by van Bevern et al. [27]:

► **Lemma 9** (adapted from van Bevern et al. [27, Lemma 2.12]). *Let (G, R, w) be an instance of RPP with integral weights, let $\varepsilon > 0$, and let $\beta = \max\{w(e) \mid e \in E(G)\}$. Then in polynomial time we can compute a weight function $w' : E(G) \rightarrow \mathbb{N}_{\geq 0}$ such that*

- $\max\{w'(e) \mid e \in E(G)\} \leq 2|E(G)|/\varepsilon$,
- and for all $\alpha \geq 1$, any solution J to (G, R, w') with weight $w'(J) \leq \alpha \text{OPT}(G, R, w')$ also fulfills $w(J) \leq \alpha \text{OPT}(G, R, w) + \varepsilon\beta$, as long as J contains at most two copies of each edge.

Proof. The rounding scheme simply sets $w'(e) := \lfloor w(e) \cdot \frac{2|E(G)|}{\varepsilon \cdot \beta} \rfloor$ for each edge e of G . This yields the first condition, by definition. For the second condition, observe that for any J we have

$$\frac{\varepsilon \cdot \beta}{2|E(G)|} w'(J) \leq w(J) \leq \frac{\varepsilon \cdot \beta}{2|E(G)|} w'(J) + \frac{\varepsilon \cdot \beta}{2|E(G)|} |J| \leq \frac{\varepsilon \cdot \beta}{2|E(G)|} w'(J) + \varepsilon\beta .$$

Hence, the two weight functions are equivalent up to scaling by a constant and the addition of at most $\varepsilon\beta$. ◀

Notice that the restriction on J having at most two copies of each edge is never a problem: whenever a solution to the RPP has three or more copies of one edge, we can delete two of them to obtain a cheaper solution.

We will combine Proposition 8 and Lemma 9 to obtain an approximation for k -component RPP whose run time does not depend on OPT .

► **Corollary 10.** *There is a randomized algorithm that, for any $\varepsilon > 0$, in time $2^{\mathcal{O}(k)}(n + \frac{1}{\varepsilon})^{\mathcal{O}(1)}$ computes a solution J for any instance (G, R, w) of RPP where $(V(G), R)$ has k connected components. The computed solution J has the property that, with constant probability, $w(J) \leq \text{OPT}(G, R, w) + \varepsilon \max\{w(e) \mid e \in J^* \dot{\cup} R\}$, where J^* is some optimal solution to the instance.*

Proof. We first guess the weight β of the most expensive edge in $J^* \dot{\cup} R$, where J^* is some optimal solution. There are only $|E(G)|$ options, so the guessing generates only polynomial overhead. All edges that are more expensive than β can be removed from the instance to get some graph G' . Now we apply Lemma 9 to get an instance with weights w' bounded by $2|E(G)|/\varepsilon$ and use the exact algorithm from Proposition 8 to get a solution J to (G', R, w') in time $2^{\mathcal{O}(k)}(n + \frac{1}{\varepsilon})^{\mathcal{O}(1)}$. From Lemma 9 with $\alpha = 1$ we know that

$$w(J) \leq \text{OPT}(G, R, w) + \varepsilon\beta \leq \text{OPT}(G, R, w) + \varepsilon \max\{w(e) \mid e \in J^* \dot{\cup} R\},$$

which proves the claim. ◀

We will be using this algorithm to complete partial solutions to instances of MDMTSP. We will need only a slight modification that allows for the presence of depots as follows.

► **Definition 11** (Depot Rural Postperson Problem). *An instance (G, D, R, w) of the DEPOT RURAL POSTPERSON PROBLEM (DRPP) consists of an RPP instance (G, R, w) and some depots $D \subseteq V(G)$. A solution is a multiset $J \subseteq E(G)$ such that $(V, R \dot{\cup} J)$ is Eulerian and each non-singleton connected component of $(V, R \dot{\cup} J)$ contains at least one depot. The weight of a solution J is $w(J) = \sum_{e \in J} w(e)$. The goal is to compute an optimal solution, which is a solution of minimum weight $\text{OPT}(G, D, R, w)$.*

The depot version DRPP can be reduced to regular RPP quite easily.

► **Corollary 12.** *There is a randomized algorithm that, for any instance (G, D, R, w) of DRPP where $(V(G), R)$ has k connected components and any $\varepsilon > 0$, in time $2^{\mathcal{O}(k \log k)}(n + \frac{1}{\varepsilon})^{\mathcal{O}(1)}$ computes a solution J such that, with constant probability, $w(J) \leq (1 + \varepsilon) \text{OPT}(G, D, R, w) + \varepsilon w(R)$.*

Proof. Note first that each connected component of $(V(G), R)$ can be assumed to contain at most one depot, so $|D| \leq k$. Some optimum solution J^* induces a partition of the connected components of $(v(G), R)$ where each partition class corresponds to those components connected to some specific depot. There are at most $|D|^k \in 2^{\mathcal{O}(k \log k)}$ possible partitions, so we can try each partition, solve the regular RPP instance on each of the $|D|$ classes of the partition using the algorithm from Corollary 10, and return the best solution we found. ◀

4 Intuition for the Algorithm

The algorithm of Xu and Rodrigues [28] executes, at a very high level, the following steps:

1. Compute a minimum-weight constrained spanning forest F for (G, D, w) .
2. Guess a set X of at most $|D| - 1$ edges such that they are in F but not in some fixed optimal tour T .
3. Discard the guessed edges X from F . This leaves at most $2|D|$ connected components in $(V, F \setminus X)$. If we have guessed correctly, every subtour of T now contains an even number of odd-degree nodes. There must exist some edges A from T such that $(F \setminus X) \cup A$ is a CSF for (G, D) with $w((F \setminus X) \cup A) \leq w(T)$. The value $|A|$ is at most $|D|$, so we also guess A .


4. Since A contains only edges from T , every subtour of T still contains an even number of odd-degree nodes with respect to $(V, (F \setminus X) \cup A)$. If we compute an odd-join J for $(F \setminus X) \cup A$, we have $w(J) \leq \frac{1}{2} \text{OPT}$, so return $((F \setminus X) \cup A) \dot{\cup} J$.

Since the algorithm needs to guess $2|D|$ edges in total (in step 2), it can be implemented in time $n^{\mathcal{O}(d)}$. We modify this guessing step by considering for discarding (in step 3) only very heavy edges, and by sidestepping the guessing of A ; instead of computing first a connected structure and then a join we do this simultaneously, using the algorithm for RPP. Specifically:

- In step 2, we only consider edges that are very expensive relative to the total weight of the forest F . If the targeted edge e is not in this collection, we do not delete it but instead use it as part of the augmenting set A , doubling the edge. This also fixes parity, but requires us to relax $w((F \setminus X) \dot{\cup} A) \leq w(T)$ to $w((F \setminus X) \dot{\cup} A) \leq (1 + \varepsilon)w(T)$. The ε can be controlled by how expensive relative to F we allow these non-deleted edges to be.
- In step 3, we do not actually guess A , we merely use its existence. We instead solve an instance of DRPP with at most $2d$ connected components for which $A \dot{\cup} J$ is a solution. Using the algorithm from Corollary 12, we can compute a $(1 + \varepsilon)$ -approximation for the DRPP in time $f(d, \varepsilon) \cdot n^{\mathcal{O}(1)}$. We use the solution J' as a replacement for $A \dot{\cup} J$ knowing $w(J') \leq (1 + \varepsilon)w(A \dot{\cup} J)$. Combining inequalities for J and F gives:

$$w((F \setminus X) \dot{\cup} J') \leq (1 + \varepsilon)w(((F \setminus X) \dot{\cup} A) \dot{\cup} J) \leq (1 + \varepsilon)\frac{3}{2} \text{OPT} .$$

An illustration of the augmentation scheme can be found in Figure 3.



■ **Figure 3** Illustration of the augmenting edges explored by our algorithm. Blue solid edges represent a CSF, dashed red edges an optimal tour. Note that the tours of d_2 and d_3 have an odd number of nodes with odd degree in the CSF. Our algorithm considers two options to remedy this. We either add two copies (green, dotted) of the edge e to the optimal tour, if e is considered to be light enough. This joins the tours of d_2 and d_3 to a single tour with an even number of odd-degree nodes. If e is considered too heavy for this, we remove e from the CSF and replace it with the edge \hat{e} (black, dotted). As \hat{e} comes from the optimal tour, this keeps the cost of the CSF below OPT and it fixes the parities.

5 Towards Faster Parity Correction

In this section, we give a formal version of the algorithm described in the previous section, which we state as Algorithm 2. We prove this algorithm to be a $(3/2 + \mathcal{O}(\varepsilon))$ -approximation for metric MDMTSP in Theorem 17. We will restate and reprove some of the results of Xu and Rodrigues [28] to ensure completeness of the presentation and to integrate properly our changes to their algorithm.

To this end, let us fix some notation throughout this section. Let (G, D, w) be a metric MDMTSP instance with $D = \{d_1, \dots, d_d\}$, an optimal tour T , and the minimum-weight CSF F for (G, D, w) that was computed Algorithm 1. We denote by T_i be the connected component of T containing d_i , by F_i the subtree of F containing d_i , and U_i the set of nodes in T_i that have odd degree with respect to the edges in F . We take $U = \bigcup_i U_i$.

By minimality of F , we already know that $w(F) \leq w(T) = \text{OPT}$. To extend F to an MDMTSP tour, we try to compute a minimum-weight matching between the nodes in U . It is a standard argument from the analysis of the Christofides-Serdyukov algorithm that if $|U_i|$ is even, T_i contains two disjoint matchings for the nodes in U_i . So if every U_i has even cardinality, then any minimum-weight matching has weight at most $\frac{1}{2} \text{OPT}$. But this is not the case, since a tree F_i might contain nodes from many different tours, so the odd-degree nodes are distributed arbitrarily. To record this “misalignment” between the trees and subtours we introduce the concept of an *alignment graph*.

► **Definition 13** (Alignment Graph). *The alignment graph H for (G, F, T) is constructed as $V(H) = D$, and*

$$E(H) = \{\{d_i, d_j\} \mid \exists e \in F \text{ s.t. } |e \cap V(T_i)| = |e \cap V(T_j)| = 1\} .$$

We also define a weight function $w_H : E(H) \rightarrow \mathbb{R}_+$ as

$$w_H((d_i, d_j)) := \min\{w(e) \mid e \in F, |V(T_i) \cap e| = |V(T_j) \cap e| = 1\} .$$

In the following, we assume that H is connected, otherwise the analysis holds independently for each connected component.

Now we take D_{odd} to be the collection of depots d_i for which $|U_i|$ is odd, and A_H to be any D_{odd} -join in H . The join A_H can be used to augment the original tour T to be connected. To do this we transfer the join to the original graph to ensure that it contains a “cheap” matching. For every edge $e = \{d_i, d_j\} \in E(A_H)$, pick an edge in $\hat{e} \in E(F)$ with $w(\hat{e}) = w_H(e)$ and $|\hat{e} \cap V(T_i)| = |\hat{e} \cap V(T_j)| = 1$. Denote by A the collection of these \hat{e} . Observe that every node in $T \dot{\cup} 2A$ has even degree, and every connected component of the graph contains an even number of nodes from U . Hence, there exists a U -join J in G with $w(J) \leq \frac{1}{2}w(T \dot{\cup} 2A)$.

Now notice that, if the edges in A have weight at most $\varepsilon w(F)$, this inequality yields that Algorithm 1 already achieves a good approximation ratio, specifically

$$w(F) + w(M) \leq w(F) + w(J) \leq (1 + \varepsilon)w(F) + \frac{1}{2}w(T) \leq \left(\frac{3}{2} + \varepsilon\right)w(T) .$$

Based on this observation, we are willing to augment T with low-weight edges from F to find a low-weight matching. Therefore, we need to distinguish between *heavy* and *light* edges.

► **Definition 14.** *Let $\varepsilon > 0$. An edge $e \in F$ is ε -light if $w(e) \leq \frac{\varepsilon}{d} \cdot w(F)$; else, it is ε -heavy.*

We now try to replace the ε -heavy edges in A with some other edges from T .

■ **Algorithm 2** Algorithm EXTENDED MULTI-DEPOT CHRISTOFIDES-SERDYUKOV.

Input: A metric MDMTSP instance (G, D, w) and a parameter $\varepsilon > 0$.
Output: A tour T such that $w(T) \leq (3/2 + \varepsilon) \text{OPT}$ with constant probability.

- 1 Compute a minimum-weight CSF F for (G, D, w) ;
- 2 Let T be the currently best MDMTSP solution, initially $2F$;
- 3 Let Y be the set of edges in F which are ε -heavy;
- 4 **foreach** $X \subseteq Y$, $|X| \leq |D|$ **do**
- 5 $F' := F \setminus X$;
- 6 Compute a solution M for the DRPP instance (G, D, F', w) using Corollary 12 ;
- 7 **if** $w(M \dot{\cup} F') < w(T)$ **then**
- 8 set $T = M \dot{\cup} F'$;
- 9 **end**
- 10 **end**

Result: T

► **Lemma 15** (compare [28, Section 2]). *Let $X \subseteq A$. Then there exist a set $\hat{A} \subseteq E(T)$ of edges such that $(F \setminus X) \cup \hat{A}$ is a CSF for (G, D) with $w((F \setminus X) \cup \hat{A}) \leq w(T)$.*

Proof. Consider the forest F' obtained from T by removing exactly one edge from each subtour. F' contains only edges from T , so it is disjoint from X . By a standard matroid exchange argument, for each $e \in X$ there is a $\hat{e} \in F'$ such that $F - e + \hat{e}$ is a CSF and $w(e) \leq w(\hat{e})$. This process can then be iterated to remove all of X . The collection of these \hat{e} is \hat{A} , giving $w((F \setminus X) \cup \hat{A}) \leq w(F') \leq w(T)$. ◀

This process of replacing augmenting edges from A with edges out of T also fulfills the key goal of putting an even number of odd-degree nodes into every connected component of some augmented MDMTSP solution. Consider the following lemma, which is in substance a version of a statement by Xu and Rodrigues [28, Theorem 2].

► **Lemma 16.** *Let A, X, \hat{A} be as in Lemma 15. Then every connected component of $T \cup (A \setminus X)$ contains an even number of nodes that have odd degree in $(F \setminus X) \cup \hat{A}$.*

Proof. Notice first that the connected components of $T \cup (A \setminus X)$ are the union of some of the subtours of T . Since the edges in \hat{A} belong to some tour, adding them to $F \setminus X$ flips the parity of the degrees of two nodes on the same tour, so the total parity of odd-degree nodes on that tour does not change. We can therefore restrict ourselves to considering the odd-degree nodes with respect to $F \setminus X$.

Recall that originally A was constructed from a D_{odd} -join A_H in the alignment graph H . So X corresponds to some edge set $X_H \subseteq A_H$, and we know that $A_H \setminus X_H$ constitutes an $(D_{\text{odd}} \Delta \bigcup_{e \in X_H} e)$ -join, where Δ denotes the symmetric difference. For multisets, the symmetric difference of some sets contains an item if and only if it is contained an odd number of times in their disjoint union. At the same time, removing an edge $e \in X$ from F with corresponding $\{d_i, d_j\} \in X_H$ changes the degree of one node in $V(T_i)$ and one node on $V(T_j)$. So, the depots whose tours contain an odd number of odd-degree nodes with respect to $F \setminus X$ are precisely $(D_{\text{odd}} \Delta \bigcup_{e \in X_H} e)$, so $A_H \setminus X_H$ joins them correctly. ◀

We are now ready to prove that Algorithm 2 returns a $(3/2 + \mathcal{O}(\varepsilon))$ -approximation, with constant probability.

► **Theorem 17.** *The tour returned by Algorithm 2 has weight at most $(3/2 + \mathcal{O}(\varepsilon)) \text{OPT}$, with constant probability.*

39:12 A $(3/2 + \varepsilon)$ -Approximation for Multiple TSP with a Variable Number of Depots

Proof. Set T' to be the tour returned by the algorithm. Now let A be the augmenting edge set for some optimal tour as before and X the set of ε -heavy edges in A . We look at the iteration of the algorithm where that X is considered for removal. From Lemma 15 we know that there exists some edge set \hat{A} with $w(F \setminus X \cup \hat{A}) \leq \text{OPT}$ and Lemma 16 implies that there is an edge set J such that $F \setminus X \dot{\cup} \hat{A} \dot{\cup} J$ is Eulerian, contains a depot in each connected component, and $w(J) \leq \frac{1}{2}w(T \cup (A \setminus X))$. Therefore, $\hat{A} \dot{\cup} J$ is a solution to the DRPP instance (G, D, F', w) . Hence, the M computed in the algorithm fulfills $w(M) \leq (1 + \varepsilon)w(\hat{A} \dot{\cup} J) + \varepsilon w(F')$. Putting all these inequalities together yields

$$\begin{aligned} w(T') &= w(M) + w(F \setminus X) \\ &\leq (w(F) - w(X) + w(\hat{A})) + \left(\frac{1}{2}(w(T) + w(A \setminus X)) + \varepsilon(w(\hat{A}) + w(J) + w(F))\right) \\ &\leq w(T) + \frac{1}{2}(w(T) + d \cdot \frac{\varepsilon}{d}w(F)) + \varepsilon(w(\hat{A}) + w(J) + w(F)) \leq \frac{3}{2}w(T) + 4\varepsilon w(T), \end{aligned}$$

where we use that $A \setminus X$ contains at most d edges, and all of them are ε -light. ◀

Notice that this algorithm will give a $(3/2 + \varepsilon)$ -approximation when called with $\varepsilon/4$ as the parameter of approximation. The additional run time cost will vanish in the \mathcal{O} -notation. The probability of success for this algorithm is the same as that for the algorithm in Proposition 8. Notice that while that algorithm is called many times, we only need it to succeed for one specific choice of X . If it fails in one of the other attempts, we do not care.

It remains to analyze the run time of this algorithm. We see that Y , the set of ε -heavy edges, has size at most $\frac{d}{\varepsilon}$, so there are only $\left(\frac{d}{\varepsilon}\right)^d$ possible values for X to be tried. Note also that each loop iteration requires the approximate solution of a DRPP instance with $\mathcal{O}(d)$ components which can be done in time $2^{\mathcal{O}(d \log d)}(n + \frac{1}{\varepsilon})^{\mathcal{O}(1)}$. The total run time then is $(1/\varepsilon)^{\mathcal{O}(d \log d)} \cdot n^{\mathcal{O}(1)}$, showing Theorem 1.

6 A Deterministic 3/2-Approximation for Graphic MDMTSP

In this section we provide a deterministic 3/2-approximation for MDMTSP when the metric is the shortest-path metric of an unweighted graph. The run time of the algorithm is $2^d \cdot n^{\mathcal{O}(1)}$.

Let (G, D) be an instance of graphic MDMTSP, where G this time is the unweighted graph inducing the shortest-path metric. Note that we can assume G to be connected. This allows us to construct TSP tours that are not much more expensive than optimal solutions to MDMTSP, which re-enables the original analysis of the Christofides-Serdyukov Algorithm.

For a given optimal MDMTSP tour T , we can extend it to a TSP tour by introducing at most $2(d - 1)$ edges. To do this, contract the subtours of T , find a spanning tree in the contracted graph, and double all the edges of that tree. We then see that the solution $F \dot{\cup} M$ returned by Algorithm 1 fulfills

$$w(F \dot{\cup} M) \leq w(T) + \frac{1}{2}(w(T) + 2(d - 1)) = \frac{3}{2}w(T) + d - 1 .$$

Notice that the additive term $d - 1$ is likely to be very small, since we know $w(T) \geq n - d$. A similar argument can also be made for metrics which are continuous in the sense that the space cannot be partitioned into two very distant parts.

► **Observation 18.** *Let (G, D, w) be an integer-weighted instance of MDMTSP for which there exists a constant L such that, for all $U \subseteq V(G)$, it holds $\min\{w(u, v) \mid u \in U, v \notin U\} \leq L$. Then Algorithm 1 returns a solution T for (G, D, w) with $w(T) \leq \frac{3}{2} \text{OPT}(G, D, w) + L(d - 1)$.*

Since we know $w(T)$ to be in $\Omega(n - d)$ also in this case, Algorithm 1 gives an *asymptotic* $3/2$ -approximation for any constant d and L . We can even get rid of the additive term in the graphic case (i.e. $L = 1$) with some additional run time.

► **Observation 19.** *There is a $3/2$ -approximation algorithm for graphic MDMTSP with run time $2^d \cdot n^{\mathcal{O}(1)}$.*

Proof. Let T be some fixed optimal tour. We start by guessing the set $D' \subseteq D$ of depots whose subtours in T contain at least one edge, generating on overhead of 2^d . Then we know that T contains a CSF F' for (G, D') with weight $|E(T)| - |D'|$. As before, we connect together all subtours of the depots in D' with $|D'| - 1$ edges, and double these edges. Then the tour $F \dot{\cup} M$ returned by Algorithm 1 fulfills

$$|E(F \dot{\cup} M)| \leq |E(T)| - |D'| + \frac{1}{2}(|E(T)| + 2(|D'| - 1)) = \frac{3}{2}|E(T)| - 1,$$

and that proves the claim. ◀

For the special case where we require each depot to have a non-empty tour, we do not even have to guess the correct subset of depots in Observation 19, yielding a $3/2$ -approximation in truly polynomial time.

7 Discussion

We have shown that metric MDMTSP admits a randomized $(3/2 + \varepsilon)$ -approximation algorithm in time $(1/\varepsilon)^{\mathcal{O}(d \log d)} \cdot n^{\mathcal{O}(1)}$, filling in the gap between the best-known polynomial approximation factor, 2, and the $3/2$ -approximation of Xu and Rodrigues in time $n^{\Theta(d)}$. However, there remain a number of natural openings for improving on our result:

- Can our algorithm be derandomized? Since we rely on the algorithm of Gutin et al. [11] to solve RPP instances, this would require a derandomization of their result. However, their algorithm relies on the Schwartz-Zippel Lemma [22, 32] for which no deterministic alternatives have been found in the last 40 years.
- Can the approximation factor be improved from $3/2 + \varepsilon$ to $3/2$? We loose some approximation quality both when determining which edges to delete from the CSF, and when solving RPP. Improving the first point would require a further refinement of the tree-rearrangement technique introduced by Xu and Rodrigues [28]. For the second point, the RPP algorithm of Gutin et al. would need to be sped up to run in strongly polynomial time. Again, their algorithm relies on algebraic techniques for which derandomization appears difficult, so a major technical innovation for k -component RPP is maybe necessary.
- Does there exist some polynomial-time α -approximation algorithm for MDMTSP with $\alpha < 2$? We know from Traub et al. [26] that any α -approximation algorithm for single-salesperson TSP implies a $(\alpha + \varepsilon)$ -approximation for MDMTSP for any constant number of depots, i.e. in time $n^{\Theta(d)}$. For instances with many depots however, the problem remains intractable. It is of particular interest that two major technical tools for the classical TSP, Christofides' Algorithm and the Subtour-Elimination LP, fail to achieve better-than-2-approximations in the multi-depot regime (see Figure 1 and Figure 2). It appears that to make progress on a polynomial-time algorithm some novel structural insights would be required.

References

- 1 Tolga Bektaş. The multiple traveling salesman problem: an overview of formulations and solution procedures. *Omega*, 34(3):209–219, 2006. doi:10.1016/j.omega.2004.10.004.
- 2 Tolga Bektaş. Formulations and benders decomposition algorithms for multidepot salesman problems with load balancing. *European J. Oper. Res.*, 216(1):83–93, 2012. doi:10.1016/j.ejor.2011.07.020.
- 3 Enrique Benavent and Antonio Martínez. Multi-depot multiple TSP: a polyhedral study and computational results. *Annals Oper. Res.*, 207:7–25, 2013. doi:10.1007/s10479-011-1024-y.
- 4 Kristóf Bérczi, Matthias Mnich, and Roland Vincze. A $3/2$ -approximation for the metric many-visits path TSP. *SIAM J. Discrete Math.*, 36(4):2995–3030, 2022. doi:10.1137/22M1483414.
- 5 Kristóf Bérczi, Matthias Mnich, and Roland Vincze. Approximations for many-visits multiple traveling salesman problems. *Omega*, 116:102816, 2023. doi:10.1016/j.omega.2022.102816.
- 6 M. Burger, Z. Su, and B. De Schutter. A node current-based 2-index formulation for the fixed-destination multi-depot travelling salesman problem. *European J. Oper. Res.*, 265(2):463–477, 2018. doi:10.1016/j.ejor.2017.07.056.
- 7 Nicos Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Carnegie-Mellon University, Management Sciences Research Group, 1976. doi:10.1007/s43069-021-00101-z.
- 8 Greg N. Frederickson. Approximation algorithms for some postman problems. *J. ACM*, 26(3):538–554, 1979. doi:10.1145/322139.322150.
- 9 Greg N. Frederickson, Matthew S. Hecht, and Chul E. Kim. Approximation algorithms for some routing problems. *SIAM J. Comput.*, 7(2):178–193, 1978. doi:10.1137/0207017.
- 10 A.M. Frieze. An extension of Christofides heuristic to the k -person travelling salesman problem. *Discrete Appl. Math.*, 6(1):79–83, 1983. doi:10.1016/0166-218X(83)90102-6.
- 11 Gregory Gutin, Magnus Wahlström, and Anders Yeo. Rural postman parameterized by the number of components of required edges. *J. Comput. Syst. Sci.*, 83(1):121–131, 2017. doi:10.1016/j.jcss.2016.06.001.
- 12 Klaus Jansen. An approximation algorithm for the general routing problem. *Inf. Process. Lett.*, 41(6):333–339, 1992. doi:10.1016/0020-0190(92)90161-N.
- 13 Imdat Kara and Tolga Bektaş. Integer linear programming formulations of multiple salesman problems and its variations. *European J. Oper. Res.*, 174(3):1449–1458, 2006. doi:10.1016/j.ejor.2005.03.008.
- 14 Anna R. Karlin, Nathan Klein, and Shayan Oveis Gharan. A (slightly) improved approximation algorithm for metric TSP. In *Proc. STOC 2021*, pages 32–45, 2021. doi:10.1145/3406325.3451009.
- 15 Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Plenum, 1972.
- 16 Gilbert Laporte, Yves Nobert, and Serge Taillefer. Solving a family of multi-depot vehicle routing and location-routing problems. *Transport. Sci.*, 22(3):161–172, 1988. doi:10.1287/trsc.22.3.161.
- 17 Waqar Malik, Sivakumar Rathinam, and Swaroop Darbha. An approximation algorithm for a symmetric generalized multiple depot, multiple travelling salesman problem. *Oper. Res. Lett.*, 35(6):747–753, 2007. doi:10.1016/j.orl.2007.02.001.
- 18 Tobias Mömke and Ola Svensson. Approximating graphic TSP by matchings. In *Proc. FOCS 2011*, pages 560–569, 2011. doi:10.1109/FOCS.2011.56.
- 19 Paul Oberlin, Sivakumar Rathinam, and Swaroop Darbha. A transformation for a heterogeneous, multiple depot, multiple traveling salesman problem. In *Proc. ACC 2009*, pages 1292–1297, 2009. doi:10.1109/ACC.2009.5160666.
- 20 SN Parragh. Solving a real-world service technician routing and scheduling problem. In *Proc. TRISTAN VII*, 2010.

- 21 Sivakumar Rathinam, Raja Sengupta, and Swaroop Darbha. A resource allocation algorithm for multivehicle systems with nonholonomic constraints. *IEEE Trans. Automat. Sci. Engin.*, 4(1):98–104, 2007. doi:10.1109/TASE.2006.872110.
- 22 Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. ACM*, 27(4):701–717, 1980. doi:10.1145/322217.322225.
- 23 András Sebő and Jens Vygen. Shorter tours by nicer ears: $7/5$ -approximation for the graph-TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. *Combinatorica*, 34(5):597–629, 2014. doi:10.1007/s00493-014-2960-3.
- 24 Anatoliy I. Serdyukov. On some extremal walks in graphs. *Upravlyaemye Sistemy*, 17:76–79, 1978. (in Russian).
- 25 Kaarthik Sundar and Sivakumar Rathinam. An exact algorithm for a heterogeneous, multiple depot, multiple traveling salesman problem. In *Proc. ICUAS 2015*, pages 366–371, 2015. doi:10.1109/ICUAS.2015.7152311.
- 26 Vera Traub, Jens Vygen, and Rico Zenklusen. Reducing path TSP to TSP. In *Proc. STOC 2020*, pages 14–27, 2020. doi:10.1137/20M135594X.
- 27 René van Bevern, Till Fluschnik, and Oxana Yu. Tsidulko. On approximate data reduction for the rural postman problem: Theory and experiments. *Networks*, 76(4):485–508, 2020. doi:10.1002/net.21985.
- 28 Zhou Xu and Brian Rodrigues. A $3/2$ -approximation algorithm for the multiple TSP with a fixed number of depots. *INFORMS J. Comput.*, 27(4):636–645, 2015. doi:10.1287/ijoc.2015.0650.
- 29 Zhou Xu and Brian Rodrigues. An extension of the Christofides heuristic for the generalized multiple depot multiple traveling salesmen problem. *European J. Oper. Res.*, 257(3):735–745, 2017. doi:10.1016/j.ejor.2016.08.054.
- 30 Zhou Xu, Liang Xu, and Brian Rodrigues. An analysis of the extended Christofides heuristic for the k -depot TSP. *Oper. Res. Lett.*, 39(3):218–223, 2011. doi:10.1016/j.orl.2011.03.002.
- 31 S. Yadlapalli, W.A. Malik, S. Darbha, and M. Pachter. A lagrangian-based algorithm for a multiple depot, multiple traveling salesmen problem. *Nonlinear Analysis: Real World Applications*, 10(4):1990–1999, 2009. doi:10.1016/j.nonrwa.2008.03.014.
- 32 Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Proc. EUROSAM 1979*, volume 72 of *Lecture Notes Comput. Sci.*, pages 216–226, 1979. doi:10.1007/3-540-09519-5_73.