



Capacitated Vehicle Routing with Non-Uniform Speeds

Gørtz, Inge Li; Molinaro, Marco; Nagarajan, Viswanath; Ravi, R.

Published in:
Lecture Notes in Computer Science

Link to article, DOI:
[10.1007/978-3-642-20807-2_19](https://doi.org/10.1007/978-3-642-20807-2_19)

Publication date:
2011

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Gørtz, I. L., Molinaro, M., Nagarajan, V., & Ravi, R. (2011). Capacitated Vehicle Routing with Non-Uniform Speeds. In *Lecture Notes in Computer Science: The 15th Conference on Integer Programming and Combinatorial Optimization* (Vol. 6655/2011, pp. 235-247). Berlin: Springer. Lecture Notes in Computer Science https://doi.org/10.1007/978-3-642-20807-2_19

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

6-2011

Capacitated Vehicle Routing with Non-uniform Speeds

Inge Li Gortz

Technical University of Denmark

Marco Molinaro

Carnegie Mellon University, molinaro@cmu.edu

Viswanath Nagarajan

IBM Watson Research Center

R. Ravi

Carnegie Mellon University, ravi@cmu.edu

Follow this and additional works at: <http://repository.cmu.edu/tepper>



Part of the [Management Sciences and Quantitative Methods Commons](#)

Published In

O. Gunluk and G.J. Woeginger (Eds.): IPCO 2011, LNCS 6655, 235-247.

This Conference Proceeding is brought to you for free and open access by Research Showcase @ CMU. It has been accepted for inclusion in Tepper School of Business by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Capacitated Vehicle Routing with Non-Uniform Speeds

Inge Li Gørtz, Marco Molinaro^{*}, Viswanath Nagarajan, and R. Ravi^{*}

¹ Technical University of Denmark

² Tepper School of Business, Carnegie Mellon University

³ IBM T.J. Watson Research Center

⁴ Tepper School of Business, Carnegie Mellon University

Abstract. The *capacitated vehicle routing problem* (CVRP) [17] involves distributing (identical) items from a depot to a set of demand locations, using a single capacitated vehicle. We study a generalization of this problem to the setting of multiple vehicles having non-uniform speeds (that we call *Heterogenous CVRP*), and present a constant-factor approximation algorithm.

The technical heart of our result lies in achieving a constant approximation to the following TSP variant (called *Heterogenous TSP*). Given a metric denoting distances between vertices, a depot r containing k vehicles having respective speeds $\{\lambda_i\}_{i=1}^k$, the goal is to find a tour for each vehicle (starting and ending at r), so that every vertex is covered in some tour and the maximum completion time is minimized. This problem is precisely Heterogenous CVRP when vehicles are uncapacitated.

The presence of non-uniform speeds introduces difficulties for employing standard tour-splitting techniques. In order to get a better understanding of this technique in our context, we appeal to ideas from the 2-approximation for scheduling in parallel machine of Lenstra et al. [15]. This motivates the introduction of a new approximate MST construction called *Level-Prim*, which is related to *Light Approximate Shortest-path Trees* [14]. The last component of our algorithm involves partitioning the Level-Prim tree and matching the resulting parts to vehicles. This decomposition is more subtle than usual since now we need to enforce correlation between the size of the parts and their distances to the depot.

1 Introduction

The capacitated vehicle routing problem (CVRP) is an extensively studied combinatorial optimization problem (see e.g., [17] and references therein). CVRP is defined on a metric space (V, d) , where V is a finite set of locations/vertices and $d : V \times V \rightarrow \mathbb{R}_+$ a distance function that is symmetric and satisfies triangle inequality. There is a depot vertex $r \in V$ that contains an infinite supply of an identical item, and each vertex $u \in V$ demands some units q_u of this item. A single vehicle of capacity $Q \geq 0$ is used to distribute the items. The objective is

^{*} Supported in part by NSF grant CCF-0728841.

to find a minimum length tour of the vehicle that satisfies all demands subject to the constraint that the vehicle carries at most Q units at any time.

CVRP is closely related to the Traveling Salesman Problem (TSP). It is clear that CVRP reduces to TSP in the absence of capacity constraint. More interestingly, a reverse relation is also known—essentially the best known approximation algorithm for CVRP [13] achieves a guarantee of $\rho + 1$, where ρ is the best approximation ratio for TSP.

In practice, it is natural to have a fleet of *multiple* vehicles that can run in parallel. The objective can then be to either minimize the sum of completion times of all the vehicles or to minimize the maximum completion time over all vehicles (or the makespan of the routing). Furthermore the vehicles can all be identical (same speed) or heterogeneous (have different speeds). In either case, it is not hard to see that the total completion time objective reduces to the usual CVRP on a single maximum-speed vehicle, and constant-factor approximation algorithms readily follow.

When the objective is to minimize the makespan with identical vehicles, ideas for approximating the regular CVRP problem using a tour-splitting heuristic introduced by Frederickson et al. [11] can be easily adapted to derive a constant-factor approximation algorithm (see below).

This motivates the *Heterogenous Capacitated Vehicle Routing Problem* (HetCVRP) that we consider. Here, a fleet of k vehicles with *non-uniform speeds* and uniform capacities is initially located at the depot vertex r . The objective is to satisfy the demands subject to the capacity constraints while minimizing the makespan. Our main result is a constant-factor approximation algorithm for HetCVRP.

Most of our algorithmic ideas lie in solving the special case of HetCVRP when there is no capacity constraint. This problem, which we call HetTSP, is a generalization of TSP that might be of independent interest. For most of this paper, we will focus on obtaining a constant-factor approximation for HetTSP.

1.1 Previous Techniques

Tour-splitting solutions: To illustrate the use of known techniques, we outline how to obtain a constant-factor approximation algorithm for HetTSP with uniform speeds [11]. First, notice that the union of the tours of OPT connects all vertices, and hence a minimum spanning tree (MST) has length at most $k \cdot \text{OPT}$. Then consider a MST, duplicate its edges and take an Euler tour C , which is of length $d(C) \leq 2k \cdot \text{OPT}$. Now split C into k segments of lengths at most $\frac{d(C)}{k}$ by removing edges. Finally, the tour for the i^{th} vehicle is obtained by connecting both endpoints of the i^{th} segment of C to the depot. Since twice the distance from the depot to any vertex is a lower bound on OPT, the length of each tour is at most $3 \cdot \text{OPT}$ and hence this solution is a 3-approximation. We remark that this can be extended to obtain an $O(1)$ -approximation for HetCVRP with uniform speeds (e.g., using Theorem 4 in Sect. 4).

At a very high level, this strategy has two main components: (1) Partitioning an MST into manageable-sized connected parts; (2) assigning these parts to

vehicles. This simple idea—which was already present in the 70’s—is the central piece of many heuristics and approximations for vehicle routing problems (e.g., [11, 13, ?, ?, ?, ?]). However, it is not clear how to employ this technique in the presence of vehicles with multiple speeds. This is because the two main components now need some correlation: a small part of the MST, which should be assigned to a slower vehicle, must also be relatively closer to the depot in order to be reachable by this vehicle.

Set-cover based solutions: For HetTSP with non-uniform speeds, previous approaches seem to give only a logarithmic approximation, as follows. Guess the optimal makespan OPT (within a constant factor). If each vehicle of speed s is given a length budget of $s \cdot \text{OPT}$, then the vehicles can collectively cover all vertices. Using an approximation algorithm for k -MST [12] (or the related orienteering problem [3, 5]) within a maximum-coverage framework (see e.g., [6]), we can obtain tours of length OPT that cover a constant fraction of all vertices. Repeating this coverage step until all vertices are covered gives a solution to HetTSP of makespan $O(\log n) \cdot \text{OPT}$. The intrinsic problem of this approach is that it is too general—in fact, the above algorithm also yields a logarithmic approximation even in the setting where the metric faced by each vehicle is arbitrary (instead of just scaling by its speed), and this generalization of HetTSP can be shown to be set-cover hard. It is unclear whether the approximation of this set-covering based approach can be improved for HetTSP.

1.2 Results, Techniques and Outline

We extend the tour-splitting approach described above to obtain the following result.

Theorem 1. *There are constant-factor approximation algorithms for HetTSP and HetCVRP.*

In order to obtain the approximation for HetTSP, we abstract the requirements of the two components in the tour-splitting strategy. First, we specify conditions which guarantee that a collection of r -rooted trees is “assignable”, that is, each vehicle can visit the nodes of the trees assigned to it within time $O(M)$ (Definition 1). The conditions in Definition 1 are based on the LP to obtain a 2-approximation for *scheduling in unrelated parallel machines* by Lenstra et al. [15].

Secondly, instead of partitioning an MST as in the previous section, we consider more structured spanning trees which we call **Level-Prim** trees. Consider grouping the vertices ordered according to their distance from r into levels, where the i th level includes all vertices within distance $2^i M$.⁵ The **Level-Prim** tree is simply the tree resulting from running Prim’s algorithm with the restriction that all nodes in a level are spanned before starting to pull in nodes from the next. A **Level-Prim** tree has two important properties: (i) The vertices along

⁵ Notice that given the rounding of vehicle speeds to powers of two, vertices in level i can only be served by vehicles of speed 2^i or higher given the makespan bound M .

every root-leaf path are monotonically nondecreasing in level and (ii) For every suffix of levels, the subgraph induced on it costs at most $O(1)$ times its induced MST. The first condition, which is the departing point from MSTs, greatly simplifies the decomposition procedure carried in the next step. The second property guarantees that we can decompose a Level-Prim tree into an assignable collection. These properties are formalized in Theorem 2.

The Level-Prim construction combines both MST and shortest-path distances from a root, so it is not surprising that this structure is related to *Light Approximate Shortest-Path Trees* (LAST) introduced by Khuller et al. [14]. Indeed, we use the existence of a suitably defined LAST in proving Theorem 2. We remark, however, that the properties guaranteed by LASTs are not enough for our purposes (see Sect. 3.2).

The third main component of our approximation for HetTSP is decomposing Level-Prim into an assignable collection of r -rooted trees. Roughly, we partition the edges of Level-Prim into subtrees while ensuring that each subtree consisting of vertices in levels up to i (and hence is at a distance of about $2^i M$ from the root) also has length approximately $2^i M$, and thus can be assigned to a vehicle of speed about 2^i . This partition, which relies on the two properties of Level-Prim, gives a collection of *unrooted* trees which is assignable. Due to the length of these trees, the extra distance to connect them to the root r can be charged to their edges, hence this collection can be turned into a r -rooted assignable collection.

In order to obtain an approximation to HetCVRP, we reduce this problem to approximating HetTSP in a suitably modified metric space. This new distance function encodes any additional trips to and from the root that a vehicle has to make if it runs out of capacity. The exact transformation is presented in Sect. 4.

1.3 Related Work

For the CVRP, the best known approximation ratio [13] is essentially $\rho+1$ where ρ is the best guarantee for TSP. The current best values for ρ are $\rho = \frac{3}{2}$ for general metrics [7], and $\rho = 1 + \epsilon$ (for any constant $\epsilon > 0$) for constant dimensional Euclidean metrics [1, 16]. This has been improved slightly to $1 + \rho \cdot (1 - \frac{1}{Q}) - \frac{1}{3Q^3}$ when $Q \geq 3$ [4]. Recently, Das and Mathieu [9] gave a quasi-polynomial time approximation scheme for CVRP on the Euclidean plane.

Several variants of TSP have been studied, most of which have a min-sum objective. One related problem with min-max objective is *nurse station location* [10], where the goal is to obtain a collection of trees (each rooted at a distinct depot) such that all vertices are covered and the maximum tree length is minimized. Even et al. [10] gave a 4-approximation algorithm for this problem, based on partitioning the MST and assigning to trees along the lines of Sect. 1.1; their second step, however, involves a non-trivial bipartite matching subproblem.

In proving the properties of Level-Prim, we use *Light Approximate Shortest-Path Trees* introduced by Khuller, Raghavachari and Young [14], building on the work on shallow-light trees of Awerbuch, Baratz and Peleg [2]. An (α, β) -LAST is a rooted tree that has (a) length at most β times the MST and (b) the distance from any vertex to the root (in the tree) is at most α times the

distance in the original metric. Khuller et al. [14] showed that every metric has an $(\alpha, 1 + \frac{2}{\alpha-1})$ -LAST (for any $\alpha > 1$) and this is best possible.

2 Model and Preliminaries

The input to the *Heterogenous TSP* (HetTSP) consists of a metric (V, d) denoting distances between vertices, a depot $r \in V$ and k vehicles with speeds $\{\lambda_i\}_{i=1}^k$ greater than or equal to 1. The goal is to find tours $\{\tau_i\}_{i=1}^k$ (starting and ending at r) for each vehicle so that every vertex is covered in some tour and which minimize the *maximum completion time* $\max_{i=1}^k \frac{d(\tau_i)}{\lambda_i}$.

At the loss of a factor of two in the approximation, we assume that the λ_i 's are all (non-negative integral) powers of 2. Then, for each integer $i \geq 0$ we use μ_i to denote the number of vehicles with speed 2^i . Let OPT denote the optimal value of this modified instance of HetTSP.

Let $G = (V, E)$ be the complete graph on vertices V with edge-weights corresponding to the distance function d . For any set $F \subseteq E$ of edges, we set $d(F) = \sum_{e \in F} d_e$. Given any (multi)graph H and a subset U of its vertices, $H[U]$ denotes the subgraph induced on U and H/U denotes the graph obtained by contracting vertices U to a single vertex (we retain parallel edges). Moreover, for any pair of vertices u, v in H , we use $d_H(u, v)$ to denote the length of the shortest path in H between u and v .

3 Algorithm for HetTSP

Assume that we have correctly guessed a value M such that $\frac{M}{2} \leq \text{OPT} \leq M$ (we address how to find M in the end of Sect. 3.3.) We partition the set of vertices V according to their distance to r :

$$V_0 = \{u \in V : d(r, u) \leq M\}, \text{ and}$$

$$V_i = \{u \in V : d(r, u) \in (2^{i-1}M, 2^iM]\}, \text{ for all } i \geq 1.$$

The vertices in V_i are referred to as *level i vertices*. For any $i \geq 0$, we use $V_{\leq i}$ as a shorthand for $\cup_{j=0}^i V_j$ and similarly $V_{< i} = \cup_{j=0}^{i-1} V_j = V_{\leq i-1}$.

We define the *level of an edge* $(u, v) \in E$ as the larger of the levels of u and v . For each $i \geq 0$, E_i denotes the edges in E of level i . Note that $d_e \leq 2^{i+1}M$ for all $e \in E_i$, since both end-points of e are in $V_{\leq i}$ and the triangle inequality bounds its distance by the two-hop path via the root. We use the notation $E_{\leq i} = \cup_{j=0}^i E_j$ and $E_{\geq i} = \cup_{j \geq i} E_j$.

3.1 Assignable Trees

We start by studying collections of trees that can be assigned to vehicles in a way that each vehicle takes time $O(M)$ to visit all of its assigned trees. In the following let α and β be integers.

Definition 1 (Assignable Trees). A collection of r -rooted trees $\bigcup_{i \geq 0} \mathcal{T}_i$ covering all vertices V is called (α, β) -assignable if it satisfies the following properties.

1. For each $i \geq 0$ and every $T \in \mathcal{T}_i$, $d(T) \leq \alpha 2^i M$.
2. For each $i \geq 0$, $\sum_{j \geq i} d(\mathcal{T}_j) \leq \beta M \sum_{j \geq i-1} 2^j \mu_j$.

Intuitively, the trees in \mathcal{T}_i can be assigned to vehicles with speed 2^i so as to complete in time $O(\alpha M)$. Condition (2) guarantees that the trees $\bigcup_{j \geq i} \mathcal{T}_j$ targeted by vehicles of speed 2^{i-1} and above stand a chance of being handled by them within makespan $O(\beta M)$. Interestingly, these minimal conditions are enough to eventually assign all trees in collection to vehicles while guaranteeing makespan $O((\alpha + \beta)M)$.

Lemma 1. Given an assignable collection $\bigcup_{i \geq 0} \mathcal{T}_i$ of r -rooted trees, we can obtain in polynomial time an $(4\alpha + 2\beta)$ -approximation for HetTSP.

To prove this lemma⁶, we show that condition (2) guarantees the existence of a fractional assignment of trees where each vehicle incurs load at most βM . Then using condition (1) and a result on scheduling on parallel machines [15], we round this assignment into an integral one while increasing the load on each vehicle by at most $2\alpha M$. We loose an extra factor of 2 to convert the trees into routes.

Fractional Assignment. Consider the bipartite graph H whose left side contains one node for each tree in $\bigcup_i \mathcal{T}_i$ and whose right side contains one node for each vehicle. (We identify the nodes with their respective trees/vehicles.) There is an arc between the tree $T \in \mathcal{T}_i$ and a vehicle of speed 2^j if $j \geq i - 1$.

Consider the following b -matching problem in H : for each tree $T \in \mathcal{T}_i$, we set $b(T) = d(T)$ and for each vehicle u of speed 2^j we set $b(u) = \beta 2^j M$. A (left-saturating) b -matching is one which fractionally assigns all $b(T)$ units of each tree T such that no vehicle u is assigned more than $b(u)$ units. Notice that a feasible b -matching gives a fractional assignment of trees where each vehicle incurs load at most βM .

Then our goal is to show the existence of a b -matching in H . Using a standard generalization of Hall's Theorem (e.g., see page 54 of [8]), we see that H has a feasible b -matching iff for every set V of trees, $\sum_{T \in V} b(T)$ is at most $\sum_{u \in N(V)} b(u)$, where $N(V)$ is the neighborhood of V . However, the structure of H allows us to focus only on sets V which are equal to $\bigcup_{j \geq i} \mathcal{T}_j$ for some i .⁷ Using this revised condition, H has a b -matching iff for all i , $\sum_{j \geq i} d(\mathcal{T}_j) \leq \beta M \sum_{j \geq i-1} 2^j \mu_j$. Since

⁶ We remark that a direct proof of Lemma 1 is also possible, but the route we take reveals more properties of the requirement at hand and could potentially be useful in tackling generalizations of HetTSP.

⁷ To see that all other inequalities are dominated by those coming from such sets, first notice that if V contains a tree in \mathcal{T}_i then $N(V)$ already contains all vehicles of speed 2^j for $j \geq i - 1$. Then adding to V extra trees in $\bigcup_{j \geq i} \mathcal{T}_j$ does not change its neighborhood and thus leads to a dominating inequality.

this is exactly condition (2) in Definition 1, it follows that H indeed has a b -matching (which can be obtained in polynomial time using any maximum flow algorithm [8]).

Scheduling Parallel Machines. We show how to round the fractional assignment obtained in the previous section. We consider each tree as a “job” and each vehicle as a “machine”, where the “processing time” $p_{T,u}$ of a tree T in a vehicle u of speed 2^j is $d(T)/2^j$; then the “makespan” of a vehicle is exactly equal to the sum of the processing times of the trees assigned to it.

Let $x_{T,u}$ denote the fraction of tree T assigned to vehicle u given by scaling down a b -matching in H (i.e., if the matching assigns d units of T to vehicle u , we have $x_{T,u} = d/d(T)$). The feasibility of the matching gives $\sum_T x_{T,u} p_{T,u} \leq \beta M$ for all u . Moreover, by construction of the edges of H , $x_{T,u} > 0$ for $T \in \mathcal{T}_i$ implies that u has speed at least 2^{i-1} . Then using property (1) of assignable trees we get that $x_{T,u} > 0$ implies $p_{T,u} \leq 2\alpha M$. These two properties guarantee that x is a feasible solution for the natural LP formulation for the scheduling problem with a feasible makespan value of βM and the maximum processing time t set to $2\alpha M$. Theorem 1 of [15] then asserts that x can be rounded into an integral assignment of trees to vehicles such that the load on any vehicle is at most $(2\alpha + \beta)M$.

As in Sect. 1.1, we can transform each tree in $\bigcup_{i \geq 0} \mathcal{T}_i$ into a cycle while at most doubling its length, which then gives a $(4\alpha + 2\beta)$ approximation for HetTSP. This concludes the proof of Lemma 1.

3.2 Level-Prim

In order to obtain an assignable collection of r -rooted trees for our instance, we formally introduce Level-Prim trees. These are the trees obtained by the the following procedure.

Algorithm 1 Level-Prim(G)

- 1: For each $i \geq 0$, let H_i be an MST for $G[V_{\leq i}/V_{< i}]$.
 - 2: **return** $\mathcal{H} = \bigcup_{i \geq 0} H_i$.
-

Note that Level-Prim trees can alternately be defined by modifying Prim’s algorithm such that nodes in level i are only considered to be added to the tree after all nodes in levels below i have already been added.

Theorem 2. *A Level-Prim tree $\mathcal{H} = \{H_i\}_{i \geq 0}$ satisfies the following:*

- *The vertex-levels along every root-leaf path are non-decreasing.*
- *For each $i \geq 0$, $\sum_{j \geq i} d(H_j) \leq 8 \cdot \text{MST}(G/V_{< i})$.*

Note that the second property in Theorem 2 mirrors the second property in Definition 1. A formal connection between the two is established via the following

lemma that uses an optimal vehicle routing solution to derive a feasible spanning tree connecting a suffix of the level sets.

Lemma 2 (Lower Bound). *For each level $\ell \geq 0$, $\text{MST}(G/V_{<\ell}) \leq M \cdot \sum_{j \geq \ell-1} 2^j \mu_j$.*

Proof. Consider an optimal solution for HetTSP and let E^* be the set of edges traversed by vehicles in this solution; label each edge in E^* by the vehicle that traversed it. Clearly E^* connects all vertices to the root r .

Only vehicles having speed at least $2^{\ell-1}$ can reach any vertex in $V_{\geq \ell}$ (since a vehicle of speed s travels distance at most $s \cdot \text{OPT} \leq s \cdot M$). Thus every edge in $E^* \cap E_{\geq \ell}$ must be labeled by some vehicle of speed at least $2^{\ell-1}$. This implies that $d(E^* \cap E_{\geq \ell}) \leq M \cdot \sum_{j \geq \ell-1} 2^j \mu_j$, since the right hand side is a bound on the total length traversed by vehicles having speed at least $2^{\ell-1}$.

On the other hand, since E^* connects all vertices, $E^* \cap E_{\geq \ell}$ contains a spanning tree of $G/V_{<\ell}$. Thus $\text{MST}(G/V_{<\ell}) \leq d(E^* \cap E_{\geq \ell}) \leq M \cdot \sum_{j \geq \ell-1} 2^j \mu_j$.

We get the following corollary of Theorem 2.

Corollary 1. *A Level-Prim tree $\mathcal{H} = \{H_i\}_{i \geq 0}$ satisfies the following:*

- *The vertex-levels along every root-leaf path are non-decreasing.*
- *For each $i \geq 0$, $\sum_{j \geq i} d(H_j) \leq 8M \sum_{j \geq i-1} 2^j \mu_j$.*

In the rest of this section, we prove Theorem 2. It is easy to see that for every ℓ , $\bigcup_{j=1}^{\ell} H_j$ spans $G[V_{\leq \ell}]$, hence the procedure produces a spanning tree for G . Moreover, by construction we obtain that every root-leaf path in \mathcal{H} traverses the levels in non-decreasing order as desired. Thus, we focus on proving the second property in the theorem.

Instead of comparing the length of the edges in \mathcal{H} with an MST, it turns out to be much easier to use a specific LAST tree as proxy for the latter. The following LAST is implicit in the construction given in [14]. Recall that a *spider* is a tree with at most one vertex (the center) having degree greater than two.

Theorem 3 ([14]). *Given any metric (V, d) with root r , there exists a spanning spider L with center r such that:*

- *For each $u \in V$, the distance from r to u in L is at most $2 \cdot d(r, u)$.*
- *The length of L is at most four times the MST in (V, d) , i.e. $d(L) \leq 4 \cdot \text{MST}$.*

We remark that we cannot use a LAST directly instead of Level-Prim since the former does not need to have the properties asserted by Theorem 2; it is easy to find a LAST which does not satisfy the first property, while Figure 1 shows that the second can also be violated by an arbitrary amount. Using these spider LASTs we can obtain the main lemma needed to complete the proof of Theorem 2.

Lemma 3. *For any graph G and any Level-Prim tree \mathcal{H} on G , we have $d(\mathcal{H}) \leq 8 \cdot \text{MST}(G)$.*

Proof. Consider a spider LAST L for G and let \mathcal{P} denote the set of all root-leaf paths in L ; note that \mathcal{P} is edge-disjoint.

Consider any root-leaf path $P = (r = u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_k)$ in \mathcal{P} . We claim that P crosses levels almost in an increasing order. Specifically, there does not exist a pair of nodes $u_i, u_j \in P$ with $i < j$, $u_i \in V_\ell$ and $u_j \in V_{\leq \ell-2}$. Suppose (for a contradiction) that this were the case; then we would have that

$$d_L(r, u_j) = d_L(r, u_i) + d_L(u_i, u_j) \geq d_L(r, u_i) \geq d(r, u_i) > 2^{\ell-1}M,$$

where the last inequality uses $u_i \in V_\ell$. On the other hand, $d(r, u_j) \leq 2^{\ell-2}M$ since $u_j \in V_{\leq \ell-2}$; so we obtain $d_L(r, u_j) > 2 \cdot d(r, u_j)$, which contradicts the definition of L (see Theorem 3).

Now we transform L into another spider L' which traverses levels in non-decreasing order as follows. For each root-leaf path $P = (r = u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_k)$, perform the following modification. Let $\{a_1, a_2, \dots, a_{k'}\}$ be the subsequence of P consisting of the vertices in *even numbered* levels, i.e. each $a_i \in L_{2\ell}$ for some $\ell \geq 0$. Similarly, let $\{b_1, b_2, \dots, b_{k''}\}$ be the subsequence of P consisting of the vertices in *odd numbered* levels. Define two paths $P_{\text{even}} := (r \rightarrow a_1 \rightarrow \dots \rightarrow a_{k'})$ (shortcutting P over nodes b_i 's) and $P_{\text{odd}} := (r \rightarrow b_1 \rightarrow \dots \rightarrow b_{k''})$ (shortcutting P over a_i 's). Observe that both P_{even} and P_{odd} cross levels *monotonically*: if not then there must be some $i < j$ in P with $u_i \in V_\ell$ and $u_j \in V_{\leq \ell-2}$, contrary to the previous claim. Also, by employing the triangle inequality we have that $d(P_{\text{even}}), d(P_{\text{odd}}) \leq d(P)$. Finally define the spider L' as the union of the paths $\{P_{\text{even}}, P_{\text{odd}}\}$ over all root-leaf paths $P \in \mathcal{P}$.

By construction, the vertex levels along each root-leaf path of L' are non-decreasing. Additionally $d(L') = \sum_{P \in \mathcal{P}} (d(P_{\text{even}}) + d(P_{\text{odd}})) \leq 2 \sum_{P \in \mathcal{P}} d(P) = 2 \cdot d(L) \leq 8 \cdot \text{MST}$, by Theorem 3. Now partition the edges of L' as:

$$\Delta_\ell = \begin{cases} L'[V_0] & \text{if } \ell = 0, \\ L'[V_{\leq \ell}] \setminus L'[V_{\leq \ell-1}] & \text{if } \ell \geq 1. \end{cases}$$

By the monotone property of paths in L' , it follows that $L'[V_{\leq \ell}]$ is connected for every $\ell \geq 0$. Thus Δ_ℓ is a *spanning tree* in graph $G[V_{\leq \ell}]/V_{< \ell}$. Since H_ℓ in the Level-Prim construction, is chosen to be an MST in $G[V_{\leq \ell}]/V_{< \ell}$, we obtain $d(H_\ell) \leq d(\Delta_\ell)$. So, $d(\mathcal{H}) = \sum_{\ell \geq 0} d(H_\ell) \leq \sum_{\ell \geq 0} \Delta_\ell = d(L') \leq 8 \cdot \text{MST}$. This completes the proof of the lemma.

Completing proof of Theorem 2. We now prove the second property in Theorem 2. Lemma 3 directly implies this property for $i = 0$. For any level $i > 0$ consider the graph $G' = G/V_{< i}$; observe that $\bigcup_{j \geq i} H_j$ is a Level-Prim tree for G' (due to the iterative construction of $\mathcal{H} = \bigcup_{\ell \geq 0} H_\ell$). Thus applying Lemma 3 to graph G' and its Level-Prim $\bigcup_{j \geq i} H_j$, we have $\sum_{j \geq i} d(H_j) \leq 8 \cdot \text{MST}(G') = 8 \cdot \text{MST}(G/V_{< i})$.

3.3 Decomposition Procedure

In this section we decompose a Level-Prim tree into an assignable collection $\bigcup_{i \geq 0} \mathcal{T}_i$ of r -rooted trees. Motivated by Corollary 1, the idea is to essentially break each subgraph H_i into many pieces and connect them to r in order to

form the set of trees \mathcal{T}_i . More specifically, assume for now that each connected component in H_i is large enough, i.e. has length at least $2^i M$. Then for each $i \geq 0$, break the connected components of H_i into trees of length approximately $2^i M$; add to each tree the shortest edge connecting them to r and set \mathcal{T}_i as the collection of r -rooted trees obtained. By construction we get that $\bigcup_{i \geq 0} \mathcal{T}_i$ satisfies the first property of an assignable collection. Moreover, notice that each edge added to connect a tree to the root has approximately the same length as the tree itself; this guarantees that $d(\mathcal{T}_i) \lesssim 2d(H_i)$. It then follows that the collection $\bigcup_{i \geq 0} \mathcal{T}_i$ is assignable.

Notice that it was crucial to break H_i into trees of length at least approximately 2^i . But this is problematic when H_i has a small connected component. In this case we show that such a small component is always attached to (“dangling” from) a large enough component in H_{i-1} (otherwise the dangling edge to a much earlier level will already make this component heavy enough not to be small); then we simply treat the small component as an integral part of the latter.

Now we formally describe the proposed decomposition procedure.

Step 1. Let \mathcal{S}_0 contain the subtree $H_0 = \mathcal{H} \cap E_0$. For each level $i \geq 1$: partition edges $\mathcal{H} \cap E_i$ into a collection \mathcal{S}_i of (unrooted) subtrees such that each subtree contains exactly one edge from $V_{<i}$ to V_i . For any $\tau \in \mathcal{S}_i$ call the unique edge from $V_{<i}$ to V_i its *head-edge* $h(\tau)$. Note that such a partition is indeed possible since $\mathcal{H}[V_{\leq i}]/V_{<i}$ is connected.

Any subtree in \mathcal{S}_i (for $i \geq 0$) is referred to as a level i subtree. Note that head-edges are defined only for subtrees in level 1 and above.

Step 2. For each level $i \geq 0$: *mark* those $\tau \in \mathcal{S}_i$ that have $d(\tau) \geq 2^{i-3}M$. In addition, *mark* the tree H_0 in \mathcal{S}_0 . Let \mathcal{S}_i^m and \mathcal{S}_i^u denote the marked and unmarked subtrees in \mathcal{S}_i .

Step 3. For each level $i \geq 1$ and unmarked $\sigma \in \mathcal{S}_i^u$: define $\pi(\sigma)$ as the subtree in $\bigcup_{j < i} \mathcal{S}_j$ containing the other end-point of $h(\sigma)$.

Claim. For $i \geq 1$ and unmarked $\sigma \in \mathcal{S}_i^u$, $\pi(\sigma) \in \mathcal{S}_{i-1}$. Moreover, $\pi(\sigma)$ is marked.

Proof. Since σ is unmarked in level $i \geq 1$, $d(h(\sigma)) \leq d(\sigma) < 2^{i-3}M$. So the end-point v of $h(\sigma)$ in $\pi(\sigma)$ satisfies $d(r, v) \geq \frac{3}{2} \cdot 2^{i-2}M$, otherwise $d(h(\sigma)) \geq 2^{i-1}M - d(r, v) > 2^{i-3}M$. In particular $v \in V_{i-1}$ and thus $\pi(\sigma) \in \mathcal{S}_{i-1}$.

For the second part of the claim, notice that if $i = 1$ then $\pi(\sigma) = H_0$, which is always marked. So suppose $i \geq 2$. From the above, $\pi(\sigma)$ is in level $i - 1 \geq 1$ and hence contains a head-edge. This implies that $\pi(\sigma)$ contains some vertex $u \in V_{<i-1}$, namely an end-point of $h(\pi(\sigma))$. But then $d(\pi(\sigma)) \geq d(u, v) \geq d(r, v) - d(r, u) \geq 2^{i-3}M$, where we used $d(r, u) \leq 2^{i-2}M$ since $u \in V_{<i-1}$ and $d(r, v) \geq \frac{3}{2} \cdot 2^{i-2}M$ from above. Thus $\pi(\sigma)$ must be marked.

Step 4. For each level $i \geq 0$ and marked $\tau \in \mathcal{S}_i^m$: define $\text{Dangle}(\tau) = \pi^{-1}(\tau)$ as the set of all unmarked $\sigma \in \mathcal{S}_{i+1}^u$ having $\pi(\sigma) = \tau$. Clearly $d(\sigma) \leq 2^{i-2}M$ for all $\sigma \in \text{Dangle}(\tau)$.

Step 5. For each level $i \geq 0$ and marked $\tau \in \mathcal{S}_i^m$: partition the tree $\tau \cup \text{Dangle}(\tau)$ into subtrees T_1, \dots, T_q such that the first $q - 1$ trees have length in the range $[2^{i+1}M, 2^{i+2}M]$ and T_q has length at most $2^{i+2}M$. Notice that this is possible

since all edges of $\tau \cup \text{Dangle}(\tau)$ belong to $E_{\leq i+1}$ and hence have length at most $2^{i+1}M$. Finally, add the shortest edge from r to each of these new subtrees to obtain a collection $\mathcal{T}_i(\tau)$ of r -rooted trees.

Claim. For any $T \in \mathcal{T}_i(\tau)$, we have $d(T) \leq 3 \cdot 2^{i+1}M$.

Proof. Notice that every $T \in \mathcal{T}_i(\tau)$ consists of a T_j (for some $1 \leq j \leq q$) and an edge from r to a node in $V_{\leq i+1}$. Since the former has length at most $2^{i+2}M$ and the latter has length at most $2^{i+1}M$, it follows that $d(T) \leq 3 \cdot 2^{i+1}M$.

Claim. $\sum_{T \in \mathcal{T}_i(\tau)} d(T) \leq 5 \cdot [d(\tau) + d(\text{Dangle}(\tau))]$.

Proof. We break the analysis into two cases depending of q . Suppose $q = 1$, namely $\mathcal{T}_i(\tau)$ consists of a single tree T . In this case $T = \tau \cup \text{Dangle}(\tau) \cup \{e\}$, where e is an edge to r . If $i = 0$ then $d(e) = 0$ and the result holds directly. If $i > 0$ then τ has a node in $V_{< i}$ and hence $d(e) \leq 2^{i-1}M$. Because τ is marked and different than H_0 , the lower bound on its length implies that $d(e) \leq 2^{i-1}M \leq 4d(\tau) \leq 4(d(\tau) + d(\text{Dangle}(\tau)))$. The result follows by adding the length of $\tau \cup \text{Dangle}(\tau)$ to both sides.

Now suppose $q > 1$. Since all trees in $\mathcal{T}_i(\tau)$ lie in $V_{\leq i+1}$, each edge from the root in $\mathcal{T}_i(\tau)$ has length at most $2^{i+1}M$. So the left hand side is at most $\sum_{j=1}^q d(T_j) + q \cdot 2^{i+1}M$. But for $j < q$ we have $d(T_j) \geq 2^{i+1}M$, so the last term of the previous expression can be upper bounded by $\frac{q}{(q-1)} \sum_{j=1}^{q-1} d(T_j)$. This bound is smallest when $q = 2$, which then gives $\sum_{T \in \mathcal{T}_i(\tau)} d(T) \leq 3 \sum_{j=1}^q d(T_j) \leq 3d(\tau)$. This concludes the proof of the claim.

Step 6. For each level $i \geq 0$: define $\mathcal{T}_i = \bigcup_{\tau \in \mathcal{S}_i^m} \mathcal{T}_i(\tau)$.

The following lemma summarizes the main property of our decomposition procedure.

Lemma 4. *The collection $\{\mathcal{T}_i\}_{i \geq 0}$ obtained from the above procedure is (6, 40)-assignable.*

Proof. By Claim 3.3, each tree in \mathcal{T}_i has length at most $3 \cdot 2^{i+1}M$. So the collection satisfies condition (1) of Definition 1.

Fix any $i \geq 0$ for condition (2) in Definition 1. Due to Corollary 1, it suffices to prove that $\sum_{j \geq i} d(\mathcal{T}_j) \leq 5 \cdot \sum_{j \geq i} d(H_j)$. Using Claim 3.3 we obtain that

$$d(\mathcal{T}_j) = \sum_{\tau \in \mathcal{S}_j^m} d(\mathcal{T}_i(\tau)) \leq 5 \cdot \sum_{\tau \in \mathcal{S}_j^m} [d(\tau) + d(\text{Dangle}(\tau))] = 5 \cdot d(\mathcal{S}_j^m) + 5 \cdot d(\mathcal{S}_{j+1}^u).$$

The last equality above uses the fact that that $\{\text{Dangle}(\tau) : \tau \in \mathcal{S}_j^m\}$ is a partition of \mathcal{S}_{j+1}^u . Thus:

$$\sum_{j \geq i} d(\mathcal{T}_j) \leq 5 \cdot \sum_{j \geq i} d(\mathcal{S}_j^m) + 5 \cdot \sum_{j \geq i} d(\mathcal{S}_{j+1}^u) \leq 5 \cdot \sum_{j \geq i} d(\mathcal{S}_j) = 5 \cdot \sum_{j \geq i} d(H_j). \quad (1)$$

This concludes the proof of Lemma 4.

Summary of the Algorithm. Our algorithm starts with an initial low guess of M and runs the Level-Prim procedure. If the second condition in Corollary 1 does not hold for this run, we double the guess for M and repeat until it is satisfied (this happens the first time that M reaches the condition for the correct guess: $\frac{M}{2} \leq \text{OPT} \leq M$). We use the decomposition in this section summarized in Lemma 4 to obtain a (6,40)-assignable collection of trees. Using Lemma 1 on this collection gives us the desired constant approximation ratio by observing that the guess M in this step obeys $M \leq 2 \cdot \text{OPT}$.

4 Generalization for HetCVRP

The input to the *Heterogenous CVRP* (HetCVRP) consists of a metric (V, d) denoting distances between vertices, depot $r \in V$ (containing an infinite supply of items), demands $\{q_v\}_{v \in V}$ and k vehicles with speeds $\{\lambda_i\}_{i=1}^k$, each having capacity Q . A solution to HetCVRP consists of tours $\{\sigma_i\}_{i=1}^k$ (starting and ending at r) for each vehicle so that all demands are satisfied and each vehicle carries at most Q items at any point in time. The objective is to minimize the *maximum completion time*, $\max_{i=1}^k \frac{d(\sigma_i)}{\lambda_i}$. We study the “split-delivery” version of CVRP here, where demand at a vertex may be served by multiple visits to it; however, our result easily extends to the “unsplit-delivery” HetCVRP.

We show that the HetCVRP problem can be reduced to HetTSP in an approximation preserving way; so we also obtain an $O(1)$ -approximation for HetCVRP. The idea in this reduction is to modify the input metric based on lower-bounds for CVRP [13]. In order to avoid ambiguity, we use OPT_{vrp} to denote the optimum for HetCVRP and OPT_{tsp} to denote the optimum for HetTSP.

Theorem 4. *Consider an instance \mathcal{I} of HetCVRP. There is a poly-time constructible instance \mathcal{J} of HetTSP such that $\text{OPT}_{\text{tsp}}(\mathcal{J}) = O(1) \cdot \text{OPT}_{\text{vrp}}(\mathcal{I})$. Moreover, a solution to \mathcal{J} of makespan M can be converted in poly-time to a solution to \mathcal{I} with makespan $O(M)$.*

Proof. Let \mathcal{I} be an instance of HetCVRP as specified above. Standard scaling arguments can be used to ensure that Q is polynomial in n and $q_v \in \{0, 1, \dots, Q\}$ for all $v \in V$ (details in the full version).

Let $G = (V, E)$ denote the complete graph on vertices V with edge-weights equal to distances d . Augment G to a new graph H by adding vertices $V' = \{v_p : v \in V, p \in [q_v]\}$, and edges $E' = \{(v, v_p) : v \in V, p \in [q_v]\}$; each edge (v, v_p) has weight $\frac{d(r, v)}{Q}$. For any vertex $v \in V$, the vertices $\{v_p : p \in [q_v]\}$ are referred to as copies of v . Let (V', ℓ) denote the metric induced on vertices V' where ℓ denotes the shortest-path distances in graph H . We let \mathcal{J} be the instance of HetTSP on metric (V', ℓ) with depot r and k vehicles having speeds $\{\lambda_i\}_{i=1}^k$. Since $Q \leq \text{poly}(n)$ this reduction runs in polynomial time.

For any graph L and subset S of vertices, let $\text{MinSt}_L(S)$ denote the minimum length Steiner tree connecting S . For any subset $T \subseteq V'$ and $v \in V$ let $N_v(T)$ denote the number of v -copies in T ; also define $\pi(T) = \{v \in V :$

$N_v(T) > 0\}$. Observe that for any $T \subseteq V'$ we have $\text{MinSt}_H(T) = \text{MinSt}_G(\pi(T)) + \sum_{v \in V} \frac{N_v(T) \cdot d(r,v)}{Q}$ by the definition of graph H .

We first show that $\text{OPT}_{\text{tsp}}(\mathcal{J}) = O(\text{OPT}_{\text{vrp}}(\mathcal{I}))$. Consider an optimal solution $\{\sigma_i\}_{i=1}^k$ to \mathcal{I} . For each $i \in [k]$, let $c_i(v) \in \{0, 1, \dots, q_v\}$ denote the units of demand at vertex $v \in V$ served by vehicle i , and let $S_i = \{v \in V : c_i(v) > 0\}$. Note that $\sum_{i=1}^k c_i(v) = q_v$ for all $v \in V$; hence we can choose $S'_i \subseteq V'$ for each $i \in [k]$ such that $\cup_{i=1}^k S'_i = V'$ and $N_v(S'_i) = c_i(v)$ for all $v \in V, i \in [k]$. Since σ_i is a capacitated tour in G serving demands $\{c_i(v) : v \in S_i\}$, we have $d(\sigma_i) \geq \max\left\{\text{MinSt}_G(\{r\} \cup S_i), \sum_{v \in S_i} \frac{c_i(v) \cdot d(r,v)}{Q}\right\}$ using the (connectivity and capacitated routing) lower-bounds for CVRP [13]. Thus $\text{MinSt}_H(\{r\} \cup S'_i) = \text{MinSt}_G(\{r\} \cup S_i) + \sum_{v \in S_i} \frac{c_i(v) \cdot d(r,v)}{Q} \leq 2 \cdot d(\sigma_i)$. Now consider the solution to \mathcal{J} where the i^{th} vehicle visits vertices S'_i along the minimum TSP tour on $\{r\} \cup S'_i$, for all $i \in [k]$; the distance traversed by the i^{th} vehicle is at most $2 \cdot \text{MinSt}_H(\{r\} \cup S'_i) \leq 4 \cdot d(\sigma_i)$. So the **HetTSP** objective value of this solution is at most $\max_{i \in [k]} \frac{4 \cdot d(\sigma_i)}{\lambda_i} = 4 \cdot \text{OPT}_{\text{vrp}}(\mathcal{I})$.

Now consider a solution $\{\tau_i\}_{i=1}^k$ to \mathcal{J} with makespan M . Let $R_i \subseteq V'$ denote the vertices that are served by each vehicle $i \in [k]$. Since τ_i is a TSP tour on $\{r\} \cup R_i$, we have $d(\tau_i) \geq \text{MinSt}_H(\{r\} \cup R_i) = \text{MinSt}_G(\{r\} \cup \pi(R_i)) + \sum_{v \in V} \frac{N_v(R_i) \cdot d(r,v)}{Q}$. Now fix $i \in [k]$ and consider the instance of CVRP on vertices $\{r\} \cup \pi(R_i)$ with demands $\{N_v(R_i) : v \in \pi(R_i)\}$. As mentioned in the previous paragraph, $\max\left\{\text{MinSt}_G(\{r\} \cup \pi(R_i)), \sum_{v \in \pi(R_i)} \frac{N_v(R_i) \cdot d(r,v)}{Q}\right\}$ is a lower-bound for this instance, and the algorithm from [13] returns a solution σ_i within a $\rho = O(1)$ factor of this lower-bound. It readily follows that $\{\sigma\}_{i=1}^k$ is a feasible solution to \mathcal{I} with makespan at most $\max_{i \in [k]} \frac{\rho \cdot d(\tau_i)}{\lambda_i} = O(M)$.

We note that this algorithm returns a *non-preemptive* **HetCVRP** solution, i.e., each item once picked up at the depot stays in its vehicle until delivered to its destination. Moreover, the lower-bounds used by the **HetCVRP** algorithm also hold for the (less restrictive) *preemptive* version, where items might be left temporarily at different vertices while being moved from the depot to their final destination. Thus our algorithm also bounds the ‘‘preemption gap’’ (ratio of optimal non-preemptive to preemptive solutions) in **HetCVRP** by a constant.

5 Open Problems

One interesting open question regards the approximability of **HetTSP** and **HetCVRP** when vehicles are located in multiple different depots across the space. The current definition of an assignable collection and the definition of **Level-Prim** crucially depend on the assumption of a unique depot, hence an extension to the multi-depot case is likely to require new ideas. Another interesting direction is to consider **HetCVRP** with non-uniform capacities, where the ideas presented in Sect. 4 do not seem to generalize directly.

References

1. S. Arora. Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. *J. ACM*, 45:753–782, September 1998.
2. B. Awerbuch, A. Baratz, and D. Peleg. Cost-sensitive analysis of communication protocols. In *Proceedings of the 9th Annual Symposium on Principles of Distributed Computing*, pages 177–187, 1990.
3. A. Blum, S. Chawla, D. R. Karger, T. Lane, A. Meyerson, and M. Minkoff. Approximation algorithms for orienteering and discounted-reward tsp. *SIAM J. Comput.*, 37(2):653–670, 2007.
4. A. Bompadre, M. Dror, and J. Orlin. Probabilistic Analysis of Unit Demand Vehicle Routing Problems. *J. Appl. Probab.*, 44:259–278, 2007.
5. C. Chekuri, N. Korula, and M. Pál. Improved algorithms for orienteering and related problems. In *SODA*, pages 661–670, 2008.
6. C. Chekuri and A. Kumar. Maximum coverage problem with group budget constraints and applications. In *APPROX-RANDOM*, pages 72–83, 2004.
7. N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. *Report 388, Graduate School of Industrial Administration, CMU*, 1976.
8. W. J. Cook, W. H. Cunningham, W. R. Pulleyblank, and A. Schrijver. *Combinatorial optimization*. John Wiley & Sons, Inc., New York, NY, USA, 1998.
9. A. Das and C. Mathieu. A Quasi-polynomial Time Approximation Scheme for Euclidean Capacitated Vehicle Routing. In *SODA*, pages 390–403, 2010.
10. G. Even, N. Garg, J. Könemann, R. Ravi, and A. Sinha. Min-max tree covers of graphs. *Oper. Res. Lett.*, 32(4):309–315, 2004.
11. G. N. Frederickson, M. S. Hecht, and C. E. Kim. Approximation algorithms for some routing problems. *SIAM J. Comput.*, 7(2):178–193, 1978.
12. N. Garg. Saving an epsilon: a 2-approximation for the k-MST problem in graphs. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computing*, pages 396–402, 2005.
13. M. Haimovich and A. H. G. R. Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10(4):pp. 527–542, 1985.
14. S. Khuller, B. Raghavachari, and N. E. Young. Balancing minimum spanning trees and shortest-path trees. *Algorithmica*, 14(4):305–321, 1995.
15. J. K. Lenstra, D. B. Shmoys, and E. Tardos. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46:259–271, 1990. 10.1007/BF01585745.
16. J. S. B. Mitchell. Guillotine Subdivisions Approximate Polygonal Subdivisions: A Simple Polynomial-Time Approximation Scheme for Geometric TSP, k-MST, and Related Problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999.
17. P. Toth and D. Vigo, editors. *The vehicle routing problem*. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia, PA, USA, 2002.

A Proof of Theorem 3

Define an (α, β) -spider as follows.

Definition 2. Given a metric (V, d) with root r , an α, β -spider is a spanning spider such that

- For each $u \in V$, the distance from r to u in L is at most $\alpha \cdot d(r, u)$.
- The length of L is at most four times the MST in (V, d) , i.e. $d(L) \leq \beta \cdot \text{MST}$.

We will show that the following algorithm produces an (α, β) -spider for G .

Algorithm 2 (α, β) -Spider

- 1: Consider an MST for G and traverse it in preorder to obtain a path $S^0 = S = (r = u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_n)$ such that $d(S^0) \leq 2 \cdot \text{MST}$.
 - 2: **for** i from 1 to n **do**
 - 3: **if** $d_S(u_1, u_i) > \alpha d(u_1, u_i)$ **then**
 - 4: **add** (u_1, u_i) to S and **mark** u_i .
 - 5: **for each** marked node u_i , remove (u_{i-1}, u_i) from S .
 - 6: **return** $S^* = S$.
-

Lemma 5. The graph S^* returned by the algorithm is a spider.

Proof. Since the algorithm keeps adding edges to the path S^0 , it is clear that before Step 5 only the root u_1 and marked nodes have degree larger than 2. Moreover, each marked node has degree exactly 3. Thus, after Step 5 we have that only the root has degree larger than 2, and the lemma follows.

Lemma 6. S^* is an (α, β) -spider.

Proof. First we prove that $d_{S^*}(u_1, u_i) \leq \alpha d(u_1, u_i)$ for all i . To see this, consider S right before Step 5. It follows from Step 4 that $d_S(u_1, u_i) \leq \alpha d(u_1, u_i)$ for all i . Noticing that S^* is a shortest path tree of S from node u_1 implies the desired result.

Now we prove that S^* satisfies the second property of (α, β) -spider. Define $v_0 = u_1$ and let v_i be the i th node marked by the algorithm. It is clear that $d(S^*) \leq d(S^0) + \sum_{i=1}^k d(u_1, v_i)$; so our goal is to upper bound the last summation.

Fix a node v_i . Consider the beginning of the iteration where v_i is marked. Notice that at this point $d_S(u_1, v_i) \leq d(u_1, v_{i-1}) + d_S(v_{i-1}, v_i)$, since edge (u_1, v_{i-1}) was already added to S ; since S^0 is subgraph of S , it is also clear that the right hand side is at most $d(u_1, v_{i-1}) + d_{S^0}(v_{i-1}, v_i)$. However, since v_i was marked, we have that $d_S(u_1, v_i) > \alpha \cdot d(u_1, v_i)$, and then using the previous bounds we obtain that $\alpha \cdot d(u_1, v_i) < d(u_1, v_{i-1}) + d_{S^0}(v_{i-1}, v_i)$.

Adding the previous inequality over all v_i 's we get that $\alpha \sum_i d(u_1, v_i) < \sum_i d(u_1, v_{i-1}) + \sum_i d_{S^0}(v_{i-1}, v_i)$. Noticing that $d(u_1, v_0) = d(u_1, u_1) = 0$ and reorganizing leads to $(\alpha - 1) \sum_i d(u_1, v_i) \leq \sum_i d_{S^0}(v_{i-1}, v_i)$. Finally, notice that

$\sum_{i=1}^k d_{S^0}(v_{i-1}, v_i) \leq d(S^0)$: this follows from traversing the path S^0 and using the triangle inequality. This gives the bound $\sum_i d(u_1, v_i) \leq d(S^0)/(\alpha - 1)$.

Plugging this back to a previous bound on the length of S^* gives $d(S^*) \leq (1 + 1/(\alpha - 1))d(S^0) \leq (2 + 2/(\alpha - 1))d(S^0)$. This concludes the proof of Theorem 3.

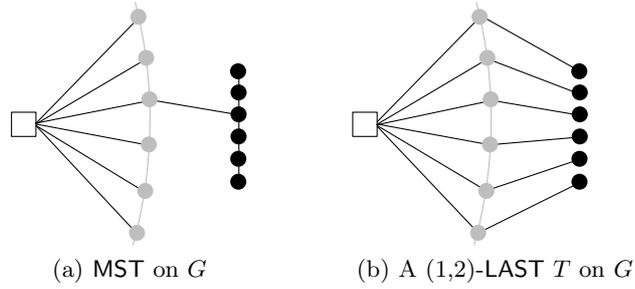


Fig. 1. Instance with depot plus n nodes in V_0 (box and gray nodes) and n nodes in V_1 (black nodes). The distance between each distinct pair of nodes in V_0 is 1, the distance from a gray node to a black node is also 1 and the distance between two distinct black nodes is $1/n$. Notice that $d(T \cap E_1) = n$ while $\text{MST}(G/V_0) < 2$.