Polynomial-time approximation schemes for k-center, k-median, and capacitated vehicle routing in bounded highway dimension

Amariah Becker

Department of Computer Science, Brown University
amariah_becker@brown.edu

Philip N. Klein

Department of Computer Science, Brown University
klein@brown.edu

David Saulpic

Département d'Informatique, École Normale Supérieure
david.saulpic@ens.fr

Abstract

The concept of bounded highway dimension was developed to capture observed properties of road networks. We show that a graph of bounded highway dimension with a distinguished *root* vertex can be embedded into a graph of bounded treewidth in such a way that u-to-v distance is preserved up to an additive error of ϵ times the u-to-root plus v-to-root distances. We show that this embedding yields a PTAS for Bounded-Capacity Vehicle Routing in graphs of bounded highway dimension. In this problem, the input specifies a depot and a set of clients, each with a location and demand; the output is a set of depot-to-depot tours, where each client is visited by some tour and each tour covers at most Q units of client demand. Our PTAS can be extended to handle penalties for unvisited clients.

We extend this embedding result to handle a set S of root vertices. This result implies a PTAS for Multiple Depot Bounded-Capacity Vehicle Routing: the tours can go from one depot to another. The embedding result also implies that, for fixed k, there is a PTAS for k-Center in graphs of bounded highway dimension. In this problem, the goal is to minimize d so that there exist k vertices (the *centers*) such that every vertex is within distance d of some center. Similarly, for fixed k, there is a PTAS for k-Median in graphs of bounded highway dimension. In this problem, the goal is to minimize the sum of distances to the k centers.

2012 ACM Subject Classification Theory of computation → Routing and network design problems

Keywords and phrases Highway Dimension, Capacitated Vehicle Routing, Graph Embeddings

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.8

Related Version A full version of this article can be found at [13], arxiv.org/abs/1707.08270.

Funding Research supported by National Science Foundation grant CCF-1409520.

Acknowledgements Thanks to Andreas Feldmann and Vincent Cohen-Addad for helpful discussions and comments.

26th Annual European Symposium on Algorithms (ESA 2018).
Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 8; pp. 8:1–8:15
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The notion of *highway dimension* was introduced by Abraham et al. [3, 1] to explain the efficiency of some shortest-path heuristics. The motivation of this parameter comes from the work of Bast et al. [11, 12] who observed that, on a road network, a shortest path from a compact region to points that are far enough must go through one of a small number of nodes. They experimentally showed that the US road network has this property, and Abraham et al. [3, 1, 2] proved results on the efficiency of shortest-path heuristics on graphs with bounded highway dimension.

Though several definitions of highway dimension have been proposed, we use the one given in [20] :

▶ **Definition 1.** The *highway dimension* of a graph $G = (V, E)$ is the smallest integer η such that for every $r \in \mathbb{R}^+$ and $v \in V$, there is a set of at most η vertices in $B_v(cr)$ such that every shortest path of length at least r that has all its vertices in $B_v(cr)$ intersects this set.

$B_v(r) = \{u \in V | d(u, v) \leq r\}$ denotes here the ball with center v and radius r. This definition is chosen as it captures this property for a wider range of transportation networks than [2]. Since the latter implies low doubling dimension, it cannot, for example, represent air traffic networks, that are star-like at large airports which causes a large doubling dimension. Nevertheless, as noted in Feldman et al. [20], these networks have a low highway dimension according to the definition of this paper (see the full version for a further discussion of these definitions).

New polynomial-time approximation schemes: Abraham et al. note that "conceivably, better algorithms for other [optimization] problems can be developed and analyzed under the small highway dimension assumption." Since road networks are thought to be modeled by graphs of small highway dimension, NP-hard optimization problems that arise in road networks are natural candidates for study. Feldmann [19] and Feldmann, Fung, Könemann, and Post [20] inaugurated this line of research, giving (respectively) a constant-factor approximation algorithm for one problem and quasi-polynomial-time approximation schemes for several other problems. In this paper, we give the first *polynomial-time approximation schemes* (PTASs) for classical optimization problems in graphs of small highway dimension.

Vehicle routing: Consider Capacitated Vehicle Routing, defined as follows. An instance consists of a positive integer Q (the *capacity*), a graph with edge-lengths, a subset Z of vertices (called *clients*), a demand function $\rho : Z \to \{1, 2 \ldots, Q\}$, and a distinguished vertex, called the *depot*. A solution consists of a set of *tours*, where each tour is a walk that starts and ends at the depot, and a function that assigns each client to a tour that passes through it, such that the total client demand assigned to each tour is at most Q. (If a client v is assigned to a tour, we say that the tour *visits* v.) The objective is to minimize the sum of lengths of the tours.

We emphasize that in this version of Capacitated Vehicle Routing, client demand is *indivisible*: a client's entire demand must be covered by a single tour. For arbitrary metrics, the problem is APX-hard, even when $Q > 0$ is fixed [9]. When Q is unbounded, it is NP-hard to approximate to within a factor of 1.5 even when the metric is that of a star [21]. Since stars have highway dimension one, this hardness result holds for graphs of bounded highway dimension. We therefore require Q to be constant. To emphasize this, we sometimes refer to the problem as Bounded-Capacity Vehicle Routing.

▶ **Theorem 2.** *For any $\epsilon > 0$, $\eta > 0$ and $Q > 0$, there is a polynomial-time algorithm that, given an instance of* Bounded-Capacity Vehicle Routing *in which the capacity is Q and the graph has highway dimension η, finds a solution whose cost is at most $1 + \epsilon$ times optimum.*

The running time is bounded by a polynomial whose degree depends on ε, η, and Q. PTASs for vehicle routing were previously known only for Euclidean spaces, although a quasi-polynomial-time approximation scheme (QPTAS) was known for planar graphs (see Section 1.2).

Our approach can be modified to handle a generalization in which an instance also specifies a *penalty* for each client, to be imposed if the solution omits the client. We also give a PTAS for a more general version of the problem, Multiple-Depot Bounded-Capacity Vehicle Routing, in which there are a constant number of depots, and each tour is required only to start and end at one of the depots.

k-Center and k-Median: Given a graph, the goal in k-Center is to select a set of k vertices (the *centers*) so as to minimize the maximum distance of a vertex to the nearest center. This problem might arise, for example, in selecting locations for k firehouses. The objective in k-Median is to minimize the average vertex-to-center distance.

For k-Center, when the number k of centers is unbounded, for any $\delta > 0$, it is NP-hard [22, 28] to obtain a $(2 - \delta)$-approximation, even in the Euclidean plane under L_1 or L_∞ metrics[1], even in unweighted planar graphs [31], and even in n-vertex graphs with highway dimension $O(\log^2 n)$ [19]. We therefore consider bounded k, but even a $(2 - \epsilon)$-approximation is $W[2]$-hard for parameter k [19] in general graphs. Thus, even for bounded k, it seems necessary to consider restricted inputs. Feldmann [19] gave a polynomial-time 3/2-approximation algorithm for bounded-highway-dimension graphs, and raised the question of whether a better approximation ratio could be achieved. The following theorem answers that question (Note that the running time is bounded by a polynomial in n whose degree does *not* depend on η, k, or ε).

▶ **Theorem 3.** *There is a function $f_1(\cdot, \cdot, \cdot)$ and a constant c such that, for each of the problems* k-Center *and* k-Median*, for any $\eta > 0$, $k > 0$ and $\varepsilon > 0$, there is an algorithm running in time $f_1(\eta, k, \varepsilon)n^c$ that, given an instance in which the graph has highway dimension at most η, finds a solution whose cost is at most $1 + \varepsilon$ times optimum.*

1.1 New metric embedding results

The key to achieving the new approximation schemes is a new result on metric embeddings of bounded-highway-dimension graphs into bounded-treewidth graphs. *Treewidth* is a measure of how complicated a graph is, and many NP-hard optimization problems in graphs become polynomial-time solvable when the input is restricted to graphs of bounded treewidth. The definition is the following.

A *tree decomposition* of a graph G is a tree T_G whose nodes are *bags* of vertices that satisfy the following three criteria: every $v \in V$ appears in at least one bag, for every edge $(u, v) \in E$ there is some bag containing both u and v and for every $v \in V$, the bags containing v form a connected subtree. The *width* of T_G is the size of the largest bag minus one, and the *treewidth* of G is the minimum width among all tree decompositions of G.

[1] Approximation better than 1.822 is hard under L_2, see [18].

A metric embedding of an (undirected) guest graph G into a host graph H is a mapping $\phi(\cdot)$ from the vertices of G to the vertices of H such that, for every pair of vertices u, v in G, the $\phi(u)$-to-$\phi(v)$ distance in H resembles the u-to-v distance in G. Usually in studying metric embeddings one seeks an embedding that preserves u-to-v distance up to some factor (the *distortion*). That is, the allowed error is proportional to the original distance. In this work, the allowed error is instead proportional to the distance from a given root vertex (or a constant number of vertices).

▶ **Theorem 4.** *There is a function $f_2(\cdot,\cdot)$ such that, for every $\varepsilon > 0$, graph G of highway dimension η, and vertex s, there exists a graph H and an embedding $\phi(\cdot)$ of G into H such that*

- *H has treewidth at most $f_2(\varepsilon, \eta)$, and*
- *for all vertices u and v, $d_G(u,v) \le d_H(\phi(u), \phi(v)) \le d_G(u,v) + \varepsilon(d_G(s,u) + d_G(s,v))$.*

As we describe in greater detail in Section 5, our PTAS for Bounded-Capacity Vehicle Routing first applies Theorem 4 with s being the depot and $\epsilon' = \epsilon/c$ for a constant c to be determined, obtaining an embedding of the original graph into the bounded-treewidth graph H. The embedding induces an instance of Vehicle Routing in H. The algorithm finds an optimal solution to this instance, and converts it to a solution for the original instance. This conversion does not increase the cost of the solution. However, we need to show that the optimal solution in the original instance induces a solution in H of not too much greater cost. We do this using a lower bound due to Haimovich and Rinnoy Kan [26].

For the multiple-depot version of vehicle routing and for k-Center and k-Median, Theorem 4 does not suffice. We present a generalization in which there is a set of root vertices, and the allowed error is proportional to the minimum distance to any root vertex.

▶ **Theorem 5.** *There is a function $f_3(\cdot,\cdot,\cdot)$ such that, for every $\varepsilon > 0$, graph G of highway dimension η and set S of vertices of G, there exists a graph H and an embedding $\phi(\cdot)$ of G into H such that*

- *H has treewidth $f_3(\eta, |S|, \varepsilon)$, and*
- *for all u and v, $d_G(u,v) \le d_H(\phi(u), \phi(v)) \le (1+O(\varepsilon))d_G(u,v) + \varepsilon \min(d_G(S,u), d_G(S,v))$*

1.2 Related Work

Metric embeddings of bounded-highway-dimension graphs: Feldmann [19] and Feldmann et al. [20] inaugurated research into approximation algorithms for NP-hard problems in bounded-highway-dimension graphs. Feldmann et al. [20] gave quasi-polynomial-time approximation schemes for Traveling Salesman, Steiner Tree, and Facility Location. The key to their results is a probabilistic metric embedding of bounded-highway dimension graphs into graphs of small treewidth. The *aspect ratio* of a graph with edge-lengths is the ratio of the maximum vertex-to-vertex distance to the minimum vertex-to-vertex distance. Feldmann et al. show that, for any $\epsilon > 0$, for any graph G of highway dimension η, there is a probabilistic embedding $\phi(\cdot)$ of G of expected distortion $1 + \epsilon$ into a randomly chosen graph H whose treewidth is polylogarithmic in the aspect ratio of G (and also depends on ϵ and η). There are two obstacles to using this embedding in achieving approximation schemes:

- The distortion is achieved only in expectation. That is, for each pair u, v of vertices, the expected $\phi(u)$-to-$\phi(v)$ distance in H is at most $(1 + \epsilon)$ times the u-to-v distance in G.
- The treewidth depends on the aspect ratio of G, so is only bounded if the aspect ratio is bounded.

The first is an obstacle for problems (e.g. k-Center) where individual distances need to be bounded; this does not apply to problems such as Traveling Salesman or Vehicle Routing where the objective is a sum of lengths of paths. The second is the reason that Feldmann et al. obtain only quasi-polynomial-time approximation schemes; it seems to be an obstacle to obtaining true PTAS. Nevertheless, the techniques introduced by Feldmann et al. are at the core of our embedding results. We build heavily on their framework.

About Vehicle Routing problem, Haimovich and Rinnoy Kan [26] proved the following lower bound[2]:

▶ **Lemma 6.** *For* Capacitated Vehicle Routing *with capacity Q, and client set Z,*

$$cost(OPT) \geq \frac{2}{Q} \sum \{d(c,s) \ : \ c \in Z\}$$

Note that the Capacitated Vehicle Routing problem is a generalization of Traveling Salesman ($Q = n$, $Z = V$, and $\rho(v) = 1, \forall v$). Conversely, Haimovich and Rinnoy Kan show how to use a solution to Traveling Salesman to achieve a constant-factor approximation for Capacitated Vehicle Routing, where the constant depends on the approximation ratio for Traveling Salesman.

Since Capacitated Vehicle Routing in general graphs is APX-hard for every fixed $Q \geq 3$ [8, 9], much work has focused on the Euclidean plane. Haimovich and Rinnoy Kan [26] gave a polynomial-time approximation scheme (PTAS) for the Euclidean plane for the case when the capacity Q is constant. Asano et al. [9] showed how to improve this algorithm to get a PTAS when Q is $O(\log n / \log\log n)$. For general capacities, Das and Mathieu [17] gave a quasi-polynomial-time approximation scheme for unbounded Q. Building on this work, Adamaszek, Czumaj, and Lingas [4] gave a PTAS that for any $\epsilon > 0$ can handle Q up to $2^{\log^\delta n}$ where δ depends on ϵ.

Little is known for higher dimensions or other metrics. Kachay gave a PTAS in $\mathbb{R}^d$ that requires Q to be $O(\log^{1/d} \log n)$ [30], and Hamaguchi and Katoh [27] and Asano, Katoh, and Kawashima [7] focused on constant-factor approximation algorithms for the case where the graph is a tree and client demand is divisible. Becker, Klein and Saulpic [14] gave the first approximation scheme for a non-Euclidean metric: they describe a quasi-polynomial-time approximation scheme in planar graphs, but only when the capacity Q is polylogarithmic in the graph size. They introduce the idea of an error that depends on the distance to the depot, which we also use in the embedding presented in our work here.

For k-Median, constant-factor approximation algorithms have been found for general metric spaces [15, 32, 29, 6]. The best known approximation ratio for k-Median in general metrics is 2.675 [15], and it is NP-hard to approximate within a factor of $1 + 2/e$ [23]. For k-Median in d-dimensional Euclidean space, PTAS have been found when k is fixed (e.g. [10]) and when d is fixed (e.g. [5]) but there exists no PTAS if k and d are part of the input [25]. Recently Cohen-Addad et al. [16] gave a local search-based PTAS for k-Median in edge-weighted planar graphs, and more generally in graphs from any nontrivial minor-closed graph family.

Outline. Section 2 provides preliminary definitions and presents useful results from Feldmann et al. [20]. In Section 3 we give an initial embedding result for graphs of bounded aspect ratio. Section 4 explains the main embedding result (Theorem 4), and Section 5 describes

[2] Although their result addresses the unit-demand case, it generalizes to instances where each non-zero client demand $\rho(v)$ is at least one.

how to use this embedding to achieve a PTAS for Capacitated Vehicle Routing, proving Theorem 2. We refer the reader to the full version [13] for a discussion of highway dimension, omitted proofs, the dynamic program for vehicle routing, and a discussion of Theorem 5 and its application to multi-depot vehicle routing, k-Center, and k- Median.

2 Preliminaries

We use OPT to denote the optimum solution for an optimization problem. For minimization problems, an α-approximation algorithm returns a solution with cost at most $\alpha \cdot \text{cost}(OPT)$. An *approximation scheme* is a family of $(1+\varepsilon)$-approximation algorithms indexed by $\varepsilon > 0$. A *polynomial-time approximation scheme* (PTAS) is an approximation scheme that for each fixed ε runs in polynomial time.

For an undirected graph $G = (V, E)$, we use $d_G(u, v)$ (or $d(u, v)$ when G is unambiguous) to denote the shortest-path distance between u and v. For any vertex subsets $W \subseteq V$ and vertex $v \in V$ we let $d(v, W)$ denote $\min_{w \in W} d(v, w)$, and we let $diam(W)$ denote $\max_{u,v \in W} d(u, v)$.

An *embedding* of a graph $G = (V, E)$ is a mapping ϕ from a *guest* graph G to a *host* graph $H = (V, E_H)$. For notational simplicity, we identify the vertices of H with points of G and therefore omit ϕ.

Let $Y \subseteq X$ be a subset of elements in a metric space (X, d). Y is a *δ-covering* of X if for all $x \in X$, $d(x, Y) \leq \delta$. Y is a *β-packing* of X if for all $y_1, y_2 \in Y$ with $y_1 \neq y_2$, $d(y_1, y_2) \geq \beta$. Y is an *ε-net* if it is both an ε-covering and an ε-packing.

Shortest-Path Covers. Now we introduce a tool for dealing with bounded highway-dimension graphs. Recall that c is a constant greater than 4.

▶ **Definition 7.** For a graph G with vertex set V and $r \in R^+$, a *shortest-path cover for scale* r $\text{SPC}(r) \subseteq V$ is a set of vertices, called *hubs*, such that every shortest path of length in $(r,\ cr/2]$ contains at least one hub. Such a cover is called *locally s-sparse* for scale r if every ball of diameter cr contains at most s vertices from $\text{SPC}(r)$.

For a graph of highway dimension η, Abraham et al. [1] showed how to find a locally $O(\eta \log \eta)$-sparse shortest-path cover in polynomial time (though they show it for a different definition of highway dimension ($c = 4$), the algorithm can be straightforwardly adapted). This result allows us to use shortest-path covers instead of directly using highway dimension.

Town Decomposition. Feldmann et al.[20] observed that a shortest-path cover for scale r naturally defines a clustering of the vertices into *towns* [20]. Informally, a *town* at scale r is a subset of vertices that are close to each other and far from other towns and from the shortest-path cover for scale r. Formally, a town is defined by at least one $v \in V$ such that $d(v, \text{SPC}(r)) > 2r$ and is composed of $\{u \in V | d(u, v) \leq r\}$. The following lemma of Feldmann et al. describes key properties of towns.

▶ **Lemma 8** (Lemma 3.2 in [20]). *If T is a town at scale r, then*
1. *$diam(T) \leq r$ and*
2. *$d(T, V \setminus T) > r$*

Feldmann et al. define a recursive decomposition of the graph using the concept of towns, which we adopt for this paper. First, scale all distances so that the shortest point-to-point distance is a little more than $c/2$. Then fix a set of scales $r_i = (c/4)^i$. We say that a town

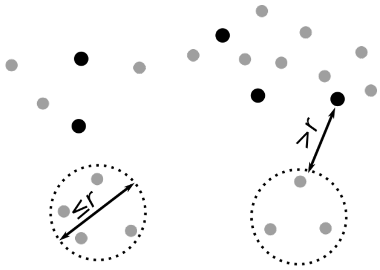

Figure 1 Illustration of Lemma 8.

T at scale r_i is on *level* i. The scaling ensures that $SPC(r_0) = \emptyset$, and therefore at level 0 every vertex forms a singleton town. The largest level is $r_{max} = \lceil \log_{c/4} diam(G_{scaled}) \rceil = \lceil \log_{c/4}(\frac{c}{2} \cdot \theta_G)) \rceil$, where θ_G is the aspect ratio of the input graph. Similarly at this topmost level, $SPC(r_{max}) = \emptyset$ since there are no shortest paths that need to be covered. The only town at scale r_{max} is the town that contains the entire graph. We say that the town at scale r_{max} and the singleton towns at scale r_0 are *trivial* towns. Since c is a constant greater than four, the total number of scales is linear in the input size.

The set $\mathcal{T} = \{T \subseteq V | T$ is a town on level $i \in \mathbb{N}\}$ of towns at all levels is called the *town decomposition.* Because of the properties of Lemma 8, this set forms a laminar family and therefore has a tree structure. Moreover, the decomposition has the following properties.

▶ **Lemma 9** (Lemma 3.3 in [20]). *For every town T in a town decomposition $\mathcal{T}$,*
1. *T has either 0 children or at least 2 children, and*
2. *if T is a town at level i and has child town T' at level j, then $j < i$.*

Approximate Core Hubs. For the purpose of approximation algorithms, it suffices to use not all hubs but a representative subset. For $\varepsilon > 0$, Feldmann et al. show how to compute, for each town T, a subset X_T of $T \bigcap \cup_i \mathrm{SPC}(r_i)$, called *approximate core hubs.* Their properties are described in Lemma 10. Recall that the *doubling dimension* of a metric is the smallest θ such that for every r, every ball of radius $2r$ can be covered by at most 2^θ balls of radius r.

▶ **Lemma 10** (Theorem 4.2 and Lemma 5.1 in [20]). *For every town $T \in \mathcal{T}$, there exist a set X_T such that:*
1. *if T_1 and T_2 are different child towns of T, and $u \in T_1$ and $v \in T_2$, then there is some $h \in X_T$ such that $d(P[u,v], h) \leq \varepsilon d(u,v)$, where $P[u,v]$ is the shortest u-to-v path, and*
2. *the doubling dimension of X_T is $\theta = O(\log(\eta s \log(1/\varepsilon))$.*

Minimality of Shortest-Path Covers. Note that the result of Lemma 10 requires the shortest-path covers be inclusion-wise minimal. For the embedding we present in Section 4, however, it is useful to assume that the depot is not a member of any town except for the trivial topmost town containing all of G and bottommost singleton town containing just the depot. This assumption can be made safely, as explained in the full version of the paper.

3 Embedding for Graphs of Bounded Aspect-Ratio

Lemma 11 describes an embedding for the case when the graph has bounded aspect-ratio, ie. the ratio between diameter and smallest distance. This embedding gives only a small *additive* error, and will prove to be a useful tool for the following sections. In this section we show how to construct this embedding.

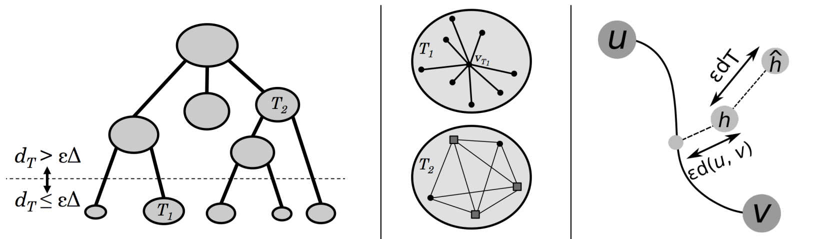

(a) Town decomposition **(b)** Embeddings **(c)** Path approximation

Figure 2 (a) An example of a town decomposition. T_1 has diameter at most $\varepsilon\Delta$ and T_2 has diameter greater than $\varepsilon\Delta$. (b) Two cases of town embeddings. T_1 is embedded as a star with center v_{T_1}. The embedding of T_2 connects all vertices in T_2 to all hubs in $\hat{X}_{T_2}$(depicted as squares). (c) Hub $\hat{h} \in \hat{X}_T$ is close to hub $h \in X_T$ which itself is close to the shortest u-to-v path.

▶ **Lemma 11.** *There is a function $f(x, y)$ such that, for any $\varepsilon > 0$ and $\eta > 0$, for any graph G with highway dimension at most η, minimal distance 1 and diameter Δ, there is a graph H with treewidth at most $f(\varepsilon, \eta)$ and an embedding $\phi(\cdot)$ of G into H such that, for all points u and v,*

$$d_G(u, v) \leq d_H(\phi(u), \phi(v)) \leq d_G(u, v) + 4\varepsilon\Delta$$

Furthermore, there is a polynomial-time algorithm to construct H and the embedding.

We first present an algorithm to compute the host graph H and a tree decomposition of H. This algorithm relies on the town decomposition $\mathcal{T}$ of G, described in Section 2.

The host graph H is constructed as follows. First, consider a town T that has diameter $d \leq \varepsilon\Delta$ but has no ancestor towns of diameter $\varepsilon\Delta$ or smaller. We call such a town a *maximal* town of diameter at most $\varepsilon\Delta$. The town T is embedded into a star: choose an arbitrary vertex v_T in T, and for each $u \in T$, include an edge in H between u and v_T with length $d_G(u, v_T)$ equal to their distance in G (see Figures 2a and 2b).

Now consider a town T of diameter $d_T > \varepsilon\Delta$. The set of approximate core hubs X_T can be used as *portals* to preserve distances between vertices lying in different child towns of T. Specifically, by Lemma 10, for every pair of vertices (u, v) in different child towns of T, X_T contains a vertex that is close to the shortest path between u and v. In order to approximate the shortest paths, it is therefore sufficient to consider a set of points *close to* X_T. Let $\hat{X}_T$ be an εd_T-net of X_T. For each $\hat{h} \in \hat{X}_T$ and $v \in T$, include an edge in H connecting v to $\hat{h}$ with length $d_H(v, \hat{h}) = d_G(v, \hat{h})$ equal to the v-to-$\hat{h}$ distance in G (see Figures 2a and 2b).

The tree decomposition D mimics the town decomposition tree: for each town T of diameter greater than $\varepsilon\Delta$, there is a bag b_T. This bag is connected in D to all of the bags of child towns of T and contains all of the vertices of the net assigned to T and of the nets assigned to T's ancestors in the town decomposition. Formally, if A_T denotes the set of all towns that contain T, $b_T = \bigcup_{T' \in A_T} \hat{X}_{T'}$. Note that if T' is the parent of T in the town decomposition, $b_T = \hat{X}_T \cup b_{T'}$. Now for each maximal town T of diameter at most $\varepsilon\Delta$ with parent town T', the tree decomposition contains a bag b_T^0 connected to a bag b_T^u for each vertex $u \in T$. We define $b_T^0 = \{v_T\} \cup b_{T'}$ and $b_T^u = \{u\} \cup b_T^0$.

Following Feldmann et al. [20], the above construction can be shown to be polynomial-time constructible. The following three lemmas therefore prove Lemma 11.

▶ **Lemma 12.** *D is a valid tree decomposition of H.*

▶ **Lemma 13.** *H has a treewidth $O((\frac{1}{\varepsilon})^{\theta}\log_{\frac{c}{4}}\frac{1}{\varepsilon})$, where θ is a bound on the doubling dimension of the sets X_T.*

Proof. Since the size of the bags is clearly bounded by the depth times the maximal cardinality of $\hat{X}_T$, it is enough to prove that, for each town T, $\hat{X}_T$ is bounded by $(\frac{1}{\varepsilon})^{\theta}$, and that the tree decomposition has a depth $O(\log_{\frac{c}{4}}\frac{1}{\varepsilon})$. By Lemma 10, the doubling dimension of X_T is bounded by θ. $\hat{X}_T$ is a subset of X_T, so its doubling dimension is bounded by 2θ (see Gupta et al. [24]). Furthermore, the aspect ratio of $\hat{X}_T$ is $\frac{1}{\varepsilon}$: the longest distance between members of $\hat{X}_T$ is bounded by the diameter d_T of the town, and the smallest distance is at least εd_T by definition of a net. The cardinality of a set with doubling dimension x and aspect ratio γ is bounded by $2^{x\lceil\log_2\gamma\rceil}$ (see [24] for a proof), therefore $|\hat{X}_T|$ is bounded by $(\frac{1}{\varepsilon})^{\theta}$. We prove now that the tree decomposition has a depth $O(\log_{\frac{c}{4}}\frac{1}{\varepsilon})$. Let T be a town of diameter $d_T > \varepsilon\Delta$ and let r_i be the scale of that town. By Lemma 8, $d_T \le r_i$, and since $r_i = (\frac{c}{4})^i$ and $d_T > \varepsilon\Delta$, we can conclude that $i > \log_{\frac{c}{4}}\varepsilon\Delta$. As the diameter of the graph is Δ, the biggest town has a diameter at most Δ. It follows that $r_i \le \Delta$ and therefore $i \le \log_{\frac{c}{4}}\Delta$. The depth of b_T in the tree decomposition is therefore bounded by $\log_{\frac{c}{4}}\frac{\Delta}{\varepsilon\Delta} = \log_{\frac{c}{4}}\frac{1}{\varepsilon}$. Furthermore, the tree decomposition of a town of diameter at most $\varepsilon\Delta$ has depth 2. The overall depth is therefore $O(\log_{\frac{c}{4}}\frac{1}{\varepsilon})$, concluding the proof. ◀

▶ **Lemma 14.** *For all vertices u and v, $d_G(u,v) \le d_H(u,v) \le d_G(u,v) + 4\varepsilon\Delta$*

Proof. Let u and v be vertices in V, and let T be the town that contains both u and v such that u and v are in different child towns of T.

If T has diameter $d_T \le \varepsilon\Delta$, then let T' be the maximal town of diameter at most $\varepsilon\Delta$ that is an ancestor of T (possibly T itself). By construction, T' was embedded into a star centered at some vertex $v_{T'} \in T'$, so $d_H(u,v) \le d_H(u,v_{T'}) + d_H(v_{T'},v) \le d_G(u,v_{T'}) + d_G(v_{T'},v) \le 2\varepsilon\Delta$.

Otherwise if T has diameter $d_T > \varepsilon\Delta$, then by Lemma 10, there is some $h \in X_T$ such that $d_G(P[u,v],h) \le \varepsilon d(u,v)$. Since $\hat{X}_T$ is an εd_T cover of X_T, there is some $\hat{h} \in \hat{X}_T$ such that $d(h,\hat{h}) \le \varepsilon d_T$. The host graph H includes edges $(u,\hat{h})$ and $(\hat{h},v)$, so

$$d_H(u,v) \le d_H(u,\hat{h}) + d_H(\hat{h},v) \le d_G(u,h) + d_G(h,v) + 2\varepsilon d(u,v) + 2\varepsilon d_T \le d_G(u,v) + 4\varepsilon\Delta$$

(see Figure 2c). Finally, since edge lengths in H are given by distances in G, $d_G(u,v) \le d_H(u,v)$ for all $u,v \in V$. ◀

4 Main Embedding: Proof of Theorem 4

4.1 Embedding Construction

Given the parameter $\hat{\epsilon}$, our goal for the embedding is that

$$d_G(u,v) \le d_H(\phi(u),\phi(v)) \le d_G(u,v) + \hat{\varepsilon}(d_G(s,u) + d_G(s,v))$$

With this goal in mind, we define $\epsilon = \min\{1/4, \hat{\varepsilon}/k\}$ for an appropriate constant k (chosen to compensate for the big-O in the following inequality), and prove that

$$d_G(u,v) \le d_H(\phi(u),\phi(v)) \le d_G(u,v) + \varepsilon(d_G(s,u) + d_G(s,v))$$

Our construction relies on the assumption that the depot s does not appear in any non-trivial town. We can make this assumption without loss of generality, as discussed in Section 2.

The root town in the composition, denoted T_0, is the town that contains the entire graph. We say that a town T that is a child of the root town is a *top-level town*, which means that the only town that properly contains T is T_0.

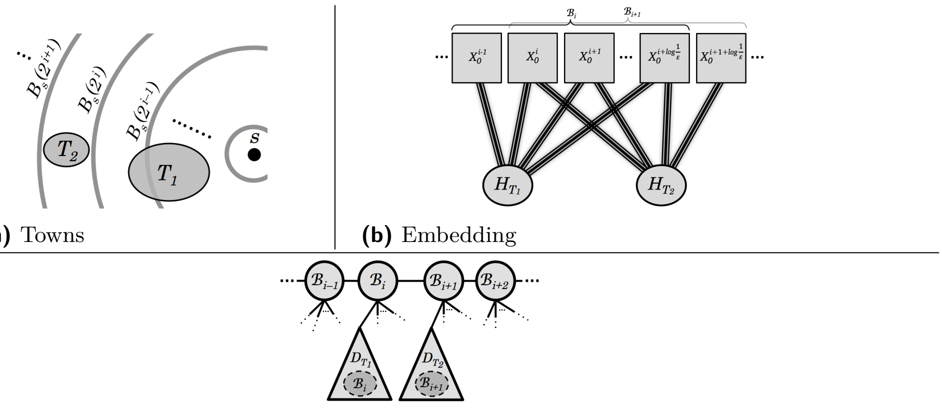

(a) Towns (b) Embedding (c) Tree decomposition

Figure 3 (a) Towns T_1 and T_2 are top-level towns, with $l(T_1) = i$ and $l(T_2) = i + 1$. (b) The embedding of each top-level town (circles) are connected to a band of $\log_2 \frac{1}{\varepsilon} + 1$ hub sets (squares). Edges are striped to convey that they connect *all* vertices of the given hub-set endpoint to *all* vertices of the town-embedding endpoint. (c) The vertices of each bag $\mathcal{B}$ (circles) are added to each bag of each descendent top-level-town tree decomposition (triangles).

The assumption that the depot, s, does not appear in any non-trivial town implies that the top-level town that contains s is the trivial singleton town. This assumption is helpful to bound the distance between a top-level town T and the depot s: as $s \notin T$, Lemma 8 gives the bound $d(T, s) \geq \text{diam}(T)$. This bound turns out to be very helpful in the construction of the host graph.

We use Lemma 11 to construct an embedding for each top-level town. It remains to connect these embeddings : we cannot approximate X_{T_0} with a net as we did in Lemma 11, because the diameter of G may be arbitrarily large.

To cope with that issue, we define inductively the hub sets $X_0^0, X_0^1, ...$ such that X_0^k is a net of $X_{T_0} \cap B_s(2^k)$. Let X_0^0 be an ε-net of $X_{T_0} \cap B_s(1)$ that contains the depot, s, and for $k \geq 0$ let X_0^{k+1} be an $\varepsilon 2^{k+1}$-net of the set $\left(X_{T_0} \cap (B_s(2^{k+1}) - B_s(2^k))\right) \bigcup X_0^k$ that contains the depot. This construction ensures that $X_0^{k+1} \cap B_s(2^k) \subseteq X_0^k$, which will be helpful in Section 4.3 to find a tree decomposition of the host graph. Note that we can assume $s \in X_{T_0}$, since adding it increases the doubling dimension by at most one and thus does not change the result of Lemma 10.

For a set of vertices $\mathcal{X} \subseteq V$, we define $l(\mathcal{X}) = \lceil \log_2(\max_{v \in \mathcal{X}} d(s, v)) \rceil$ (See Figure 3a).

For every child town T of T_0, the host graph connects every vertex v of T to every hub h in $X_0^{l(T)}, \ldots, X_0^{l(T)+\log_2(1/\varepsilon)}$ with an edge of length $d_G(v, h)$ (See Figure 3b).

4.2 Proof of Error Bound

In Lemma 16 we prove a bound on the error incurred by the embedding. Our proof makes use of the following lemma.

▶ **Lemma 15** (see full version). *For all k, X_0^k is an $\varepsilon 2^{k+1}$-covering of $X_{T_0} \cap B_s(2^k)$.*

▶ **Lemma 16.** *For all vertices u and v,*
$d_G(u, v) \leq d_H(u, v) \leq d_G(u, v) + O(\varepsilon)(d_G(s, u) + d_G(s, v))$

Proof. Consider two vertices u and v. Let T_u and T_v denote the top-level towns that contain u and v, respectively. There are two cases to consider.

(a) u and v are both connected to $\hat{h}$. **(b)** v is not connected to $\hat{h}$.

Figure 4 The shortest path between u and v in G is indicated by the curved, directed lines. The path in the host graph is represented by the straight lines.

If $T_u = T_v$, Lemma 8 gives $d_G(u,v) \leq diam(T_u) \leq d_G(T_u, V \setminus T_u)$, and therefore $diam(T_u) \leq \min\{d_G(s,u),\ d_G(s,v)\}$. Because $T_u = T_v$ is a top-level town, its embedding is given by Lemma 11, which directly gives the desired bound.

Otherwise $T_u \neq T_v$. Without loss of generality, assume that $d_G(u,s) \geq d_G(v,s)$. We show that there exists some X_0^k connected to u with a vertex $\hat{h} \in X_0^k$ close to $P[u,v]$.

By definition of the approximate core hubs, there exists $h \in X_{T_0}$ such that $d(h, P[u,v]) \leq \varepsilon d(u,v)$. Moreover, $h \in B_s(2^{l(T_u)+2})$:

$$\begin{aligned} d(s,h) &\leq d(s,u) + d(u,h) \leq d(s,u) + (1+\varepsilon)d(u,v) \\ &\leq d(s,u) + (1+\varepsilon)\left(d(s,u) + d(s,v)\right) \qquad \text{by triangle inequality} \\ &\leq d(s,u) + (1+\varepsilon)\cdot 2d(s,u) \qquad \text{since } d(u,s) \geq d(v,s) \\ &\leq (3+2\varepsilon)2^{l(T_u)} \leq 2^{l(T_u)+2} \end{aligned}$$

Since $h \in X_{T_0} \cap B_s(2^{l(T_u)+2})$, then by Lemma 15, there is an $\hat{h} \in X_0^{l(T_u)+2}$ such that $d(\hat{h}, h) \leq \varepsilon 2^{l(T_u)+3}$. Since $\log_2 \frac{1}{\varepsilon} \geq 2$, u is connected to $\hat{h}$ in the host graph.

Depending on v, there remain two cases: either v is connected to $\hat{h}$ (see Figure 4a) or not (Figure 4b). First, if v is connected to $\hat{h}$ in the host graph, $d_H(v,\hat{h}) = d_G(v,\hat{h})$ (and the same holds for u). The triangle inequality gives therefore,

$$d_H(u,v) \leq d_G(u,\hat{h}) + d_G(v,\hat{h}) \leq \underbrace{d_G(u,h) + d_G(v,h)}_{\leq (1+2\varepsilon)d_G(u,v) \text{ by definition of } h} + \underbrace{2d_G(\hat{h},h)}_{\leq 2\varepsilon 2^{l(T_u)+3} = O(\varepsilon)d(s,u)}$$

Since $d_G(u,v) \leq d_G(s,u) + d_G(s,v)$, we infer $d_H(u,v) \leq d_G(u,v) + O(\varepsilon)(d_G(s,u) + d_G(s,v))$.

Otherwise, v is not connected to $\hat{h}$. That means that either $l(T_u) + 2 < l(T_v)$ or $l(T_u) + 2 > l(T_v) + \log_2 \frac{1}{\varepsilon}$. We exclude the first case by noting that since the diameter of a town is less than its distance to the depot, $d_G(v,s) \leq d_G(u,s)$ implies that $l(T_v) \leq l(T_u) + 1$. The second case implies that $d_G(s,u) \geq O(\frac{1}{\varepsilon})d_G(s,v)$. Since the host graph connects the source s to all the vertices, $d_H(u,v) \leq d_G(s,u) + d_G(s,v) \leq d_G(u,v) + 2d_G(s,v) \leq d_G(u,v) + O(\varepsilon)(d_G(s,u) + d_G(s,v))$. ◀

4.3 Tree Decomposition

We present here the construction of a bounded-width tree decomposition D of the host graph.

For each $k > 0$ let $\mathcal{B}_k = \bigcup_{i=k-1}^{k+\log_2(1/\varepsilon)} X_0^i$. For a top-level town T, the tree decomposition D connects the decomposition D_T given by Lemma 11 to the bag $\mathcal{B}_{l(T)}$. Moreover, we add all vertices that appear in $\mathcal{B}_{l(T)}$ to all bags in the tree D_T. Finally, for every k we connect $\mathcal{B}_k$ to both $\mathcal{B}_{k-1}$ and $\mathcal{B}_{k+1}$ in D. (See Figure 3b.)

▶ **Lemma 17** (see full version)**.** *D is a valid tree decomposition of the host graph H.*

▶ **Lemma 18.** *For all k, $|X_0^k| \leq (\frac{2}{\varepsilon})^\theta$.*

Proof. Since X_0^k is a subset of X_{T_0}, it has doubling dimension 2θ (see Lemma 10). Since X_0^k is a $\varepsilon 2^k$-net, the smallest distance between two hubs in X_0^k is at least $\varepsilon 2^k$. Moreover, since $X_0^k \subseteq B_s(2^k)$, the longest distance between two hubs is at most $2 \cdot 2^k$, therefore, X_0^k has an aspect ratio of at most $\frac{2}{\varepsilon}$. The bound used in Lemma 13 on the cardinality of a set using its aspect ratio and its doubling dimension concludes the proof. ◀

▶ **Lemma 19.** *The tree decomposition D has bounded width.*

Proof. This follows from Lemma 18 together with the fact that a bag $\mathcal{B}_i$ is the union of $\log_2 \frac{1}{\varepsilon} + 2$ sets X_0^k. Lemma 13 allows to conclude. ◀

5 Capacitated Vehicle Routing

5.1 PTAS for Bounded Highway Dimension

The algorithm works as follows. The input graph G is embedded into a host graph H of bounded treewidth using the embedding given in Theorem 4. The algorithm then optimally solves the Capacitated Vehicle Routing problem with capacity Q for H, using a classical dynamic programming approach (described in the full version). The solution for H is then *lifted* to a solution in G: for each tour in the solution for H, a tour in G that visits the same clients in the same order is added to the solution for G.

We show that the embedding given in Theorem 4 is such that an optimal solution in the host graph H gives a $(1+\varepsilon)$ solution in G. Furthermore, the embedding ensures that H has small treewidth, allowing Capacitated Vehicle Routing to be solved exactly in polynomial time using dynamic programming. Putting these together gives Theorem 2.

Given an embedding with the properties described in Theorem 4, all that remains in proving Theorem 2 is showing how to solve Capacitated Vehicle Routing optimally on the host graph H and proving that such an optimal solution has a corresponding *near-optimal* solution in G. We do so in the following two lemmas (the first is proved in the full version of the paper)

▶ **Lemma 20.** *Given a graph with bounded treewidth ω and a capacity $Q > 0$,* Capacitated Vehicle Routing *can be solved optimally in $n^{O(\omega Q)}$ time.*

▶ **Lemma 21.** *For an embedding with the properties given by Theorem 4, the cost of an optimal solution in the host graph H is within a $(1 + O(\varepsilon))$-factor of the cost of the optimal solution in the guest graph G.*

Proof. Let OPT_H be the optimal solution in the host graph H and OPT_G be the optimal solution in G. A solution is described by the order in which the clients and the depot are visited: $(u, v) \in S$ indicates that the solution S visits the client v immediately after visiting u. We want to prove that $\mathrm{cost}_G(\mathrm{OPT}_H) \leq (1 + O(\varepsilon))\mathrm{cost}_G(\mathrm{OPT}_G)$.

First, since $d_G \leq d_H$, $\mathrm{cost}_G \leq \mathrm{cost}_H$. Second, the solution OPT_G is also a solution in the host graph H, since the vertices of G and H are the same. So, by definition of OPT_H, $\mathrm{cost}_H(\mathrm{OPT}_H) \leq \mathrm{cost}_H(\mathrm{OPT}_G)$. It is therefore sufficient to prove that $\mathrm{cost}_H(\mathrm{OPT}_G) \leq (1 + O(\varepsilon))\mathrm{cost}_G(\mathrm{OPT}_G)$.

By definition of cost, $\text{cost}_H(\text{OPT}_G) = \sum_{(u,v)\in\text{OPT}_G} d_H(u,v)$. Applying Theorem 4 gives

$$\text{cost}_H(\text{OPT}_G) \leq \sum_{(u,v)\in\text{OPT}_G} d_G(u,v) + O(\varepsilon)(d_G(s,u) + d_G(s,v))$$

The right side of the inequality can be rewritten as

$$\underbrace{\sum_{(u,v)\in\text{OPT}_G} d_G(u,v)}_{=\ \text{cost}_G(\text{OPT}_G)} \quad + \quad \underbrace{O(\varepsilon) \sum_{(u,v)\in\text{OPT}_G} d_G(s,u) + d_G(s,v)}_{=\ O(\varepsilon)\sum_{v\in Z} 2d_G(s,v)\ \leq\ O(\varepsilon)Q\text{cost}_G(\text{OPT}_G)\quad (*)}$$

To get the inequalities $(*)$, it is enough to remark that OPT_G visits every client exactly once and then to apply Lemma 6. As Q is constant, the whole inequality becomes

$$\text{cost}_H(\text{OPT}_G) \leq \text{cost}_G(\text{OPT}_G) + O(\varepsilon)\text{cost}_G(\text{OPT}_G) = (1 + O(\varepsilon))\text{cost}_G(\text{OPT}_G) \qquad ◀$$

References

1 Ittai Abraham, Daniel Delling, Amos Fiat, Andrew V Goldberg, and Renato F Werneck. VC-dimension and shortest path algorithms. In *International Colloquium on Automata, Languages, and Programming*, pages 690–699. Springer, 2011.

2 Ittai Abraham, Daniel Delling, Amos Fiat, Andrew V. Goldberg, and Renato F. Werneck. Highway dimension and provably efficient shortest path algorithms. *Journal of the ACM*, 63(5):41:1–41:26, 2016.

3 Ittai Abraham, Amos Fiat, Andrew V Goldberg, and Renato F Werneck. Highway dimension, shortest paths, and provably efficient algorithms. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 782–793. SIAM, 2010.

4 Anna Adamaszek, Artur Czumaj, and Andrzej Lingas. PTAS for k-tour cover problem on the plane for moderately large values of k. *International Journal of Foundations of Computer Science*, 21(6):893–904, 2010. doi:10.1142/S0129054110007623.

5 Sanjeev Arora, Prabhakar Raghavan, and Satish Rao. Approximation schemes for Euclidean k-medians and related problems. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing (STOC)*, pages 106–113, 1998.

6 Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k-median and facility location problems. *SIAM Journal on Computing*, 33(3):544–562, 2004.

7 Tetsuo Asano, Naoki Katoh, and Kazuhiro Kawashima. A new approximation algorithm for the capacitated vehicle routing problem on a tree. *Journal of Combinatorial Optimization*, 5(2):213–231, 2001.

8 Tetsuo Asano, Naoki Katoh, Hisao Tamaki, and Takeshi Tokuyama. Covering points in the plane by k-tours: a polynomial approximation scheme for fixed k. *IBM Tokyo Research Laboratory Research Report RT0162*, 1996.

9 Tetsuo Asano, Naoki Katoh, Hisao Tamaki, and Takeshi Tokuyama. Covering points in the plane by k-tours: towards a polynomial time approximation scheme for general k. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 275–283, 1997.

10 M. Bădoiu, S. Har-Peled, and P. Indyk. Approximate clustering via core-sets. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 250–257, 2002.

11 H. Bast, Stefan Funke, and Domagoj Matijevic. Ultrafast shortest-path queries via transit nodes. In Camil Demetrescu, Andrew V. Goldberg, and David S. Johnson, editors, *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, volume 74 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 175–192. American Mathematical Society, 2009.

12 H. Bast, Stefan Funke, Domagoj Matijevic, Peter Sanders, and Dominik Schultes. In transit to constant time shortest-path queries in road networks. In *Proceedings of the Ninth Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 46–59. SIAM, 2007.

13 Amariah Becker, Philip N. Klein, and David Saulpic. Polynomial-time approximation schemes for k-center and bounded-capacity vehicle routing in metrics with bounded highway dimension. *arXiv:1707.08270*, 2017.

14 Amariah Becker, Philip N Klein, and David Saulpic. A quasi-polynomial-time approximation scheme for vehicle routing on planar and bounded-genus graphs. In *25th Annual European Symposium on Algorithms (ESA)*, volume 87, pages 12:1–12:15, 2017.

15 Jaroslaw Byrka, Thomas Pensyl, Bartosz Rybicki, Aravind Srinivasan, and Khoa Trinh. An improved approximation for k-median and positive correlation in budgeted optimization. *ACM Transactions on Algorithms*, 13(2), 2017.

16 Vincent Cohen-Addad, Philip N. Klein, and Claire Mathieu. Local search yields approximation schemes for k-means and k-median in Euclidean and minor-free metrics. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2016.

17 Aparna Das and Claire Mathieu. A quasipolynomial-time approximation scheme for Euclidean capacitated vehicle routing. *Algorithmica*, 73(1):115–142, 2015.

18 Tomás Feder and Daniel Greene. Optimal algorithms for approximate clustering. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing (STOC)*, 1988.

19 Andreas Emil Feldmann. Fixed parameter approximations for k-center problems in low highway dimension graphs. In *Proceedings, Part II, of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 588–600. Springer-Verlag, 2015.

20 Andreas Emil Feldmann, Wai Shing Fung, Jochen Könemann, and Ian Post. A $(1+\varepsilon)$-embedding of low highway dimension graphs into bounded treewidth graphs. In *Proceedings, Part I, of the International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 469–480. Springer, 2015.

21 Bruce L Golden and Richard T Wong. Capacitated arc routing problems. *Networks*, 11(3):305–315, 1981.

22 Teofilo F Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 38:293–306, 1985.

23 Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *Journal of Algorithms*, 31(1):228–248, 1999.

24 Anupam Gupta, Robert Krauthgamer, and James R Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 534–543. IEEE, 2003.

25 Venkatesan Guruswami and Piotr Indyk. Embeddings and non-approximability of geometric problems. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2003.

26 Mark Haimovich and A. H. G. Rinnooy Kan. Bounds and heuristics for capacitated routing problems. *Mathematics of Operations Research*, 10(4):527–542, 1985.

27 Shin Hamaguchi and Naoki Katoh. A capacitated vehicle routing problem on a tree. In *International Symposium on Algorithms and Computation*, pages 399–407. Springer, 1998.

28 Dorit S Hochbaum and David B Shmoys. A best possible heuristic for the k-center problem. *Mathematics of Operations Research*, 10(2):180–184, 1985.

29 Kamal Jain and Vijay V. Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM*, 48(2):274–296, 2001.

30 Michael Khachay and Roman Dubinin. PTAS for the Euclidean capacitated vehicle routing problem in $\mathbb{R}^d$. In *Proceedings of the 9th International Conference on Discrete Optimization and Operations Research (DOOR)*, pages 193–205. Springer, 2016.

31 Ján Plesník. On the computational complexity of centers located in a graph. *Aplikace matematiky*, 25(6):445–452, 1980.

32 David B Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 265–274. ACM, 1997.