

Polynomial Time Approximation Schemes for Clustering in Low Highway Dimension Graphs

Andreas Emil Feldmann 

Charles University, Prague, Czech Republic
feldmann.a.e@gmail.com

David Saulpic 

LIP6, Sorbonne Université, Paris, France
david.saulpic@lip6.fr

Abstract

We study clustering problems such as k -Median, k -Means, and Facility Location in graphs of low highway dimension, which is a graph parameter modeling transportation networks. It was previously shown that approximation schemes for these problems exist, which either run in quasi-polynomial time (assuming constant highway dimension) [Feldmann et al. SICOMP 2018] or run in FPT time (parameterized by the number of clusters k , the highway dimension, and the approximation factor) [Becker et al. ESA 2018, Braverman et al. 2020]. In this paper we show that a polynomial-time approximation scheme (PTAS) exists (assuming constant highway dimension). We also show that the considered problems are NP-hard on graphs of highway dimension 1.

2012 ACM Subject Classification Theory of computation → Facility location and clustering

Keywords and phrases Approximation Scheme, Clustering, Highway Dimension

Digital Object Identifier 10.4230/LIPIcs.ESA.2020.46

Related Version A full version of the paper is available at <http://arxiv.org/abs/2006.12897>.

Funding *Andreas Emil Feldmann*: Supported by the Czech Science Foundation GAČR (grant #19-27871X), and by the Center for Foundations of Modern Computer Science (Charles Univ. project UNCE/SCI/004).

Acknowledgements We thank Vincent Cohen-Addad for helpful discussions.

1 Introduction

Clustering is a standard optimization task that seeks a “good” partition of a metric space, such that two points that are “close” should be in the same part. A good clustering of a dataset allows to retrieve and exploit data, and is therefore a common routine in data analysis. The underlying data can come from various sources and represent many different objects. In particular, it is often interesting to cluster geographic data. In that case, the metric space can be given by a transportation network, which can be modeled by graphs with low highway dimension.

In this article, we study some popular clustering objectives, namely FACILITY LOCATION, k -MEDIAN, and k -MEANS, in graphs with constant highway dimension. The two latter problems seek to find a set S of k points called *centers* in a metric (V, dist) that minimizes $\sum_{v \in V} (\min_{f \in S} \text{dist}(v, f))^p$, with $p = 1$ for k -MEDIAN and $p = 2$ for k -MEANS. The objective for FACILITY LOCATION is slightly different: each point f of the metric space has an *opening cost* w_f , and the goal is to find a set S that minimizes $\sum_{f \in S} w_f + \sum_{v \in V} \min_{f \in S} \text{dist}(v, f)$. These problems are APX-hard in general metric spaces [4].

To bypass the hardness of approximation known for these problems, researchers have considered low dimensional input, such as Euclidean spaces of fixed dimension, metrics with bounded doubling dimension, or with bounded genus. Many algorithmic tools were developed



© Andreas Emil Feldmann and David Saulpic;
licensed under Creative Commons License CC-BY
28th Annual European Symposium on Algorithms (ESA 2020).

Editors: Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders; Article No. 46; pp. 46:1–46:22



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

for that purpose: in their seminal work, Arora et al. [3] gave the first polynomial time approximation scheme (PTAS) for k -MEDIAN in \mathbb{R}^2 , which generalizes to a quasi-polynomial time approximation scheme (QPTAS) in \mathbb{R}^d for fixed d . This result was generalized by Talwar [20], who gave a QPTAS for metrics with bounded doubling dimension, and more recently by Cohen-Addad et al. [10], who gave a near-linear time approximation scheme.

In this work we focus on transportation networks, for which it can be argued that metric spaces with bounded doubling dimension are not a suitable model: for instance, hub-and-spoke networks seen in air traffic networks do not have low doubling dimension. Therefore we study graphs with constant *highway dimension*, which formalize structural properties of such networks. The following definition is taken from Feldmann et al. [14]. Here the ball $\beta_v(r)$ of radius r around $v \in V$ is the set of all vertices at distance at most r from v .

► **Definition 1.** *The highway dimension of a graph G is the smallest integer h such that, for some universal constant $c > 4$, for every $r \in \mathbb{R}^+$, and every ball $\beta_v(cr)$ of radius cr , there are at most h vertices in $\beta_v(cr)$ hitting all shortest paths of length more than r that lie in $\beta_v(cr)$.*

For this class of graphs, the only known approximation algorithms for clustering that compute $(1 + \varepsilon)$ -approximations for any $\varepsilon > 0$ either run in quasi-polynomial time, i.e., QPTASs [14], or with runtime $f(h, k, \varepsilon) \cdot n$ for some exponential function f , i.e., parameterized approximation schemes [6, 8]. Thus an open problem is to identify polynomial-time approximation schemes (PTASs) for clustering in graphs of constant highway dimension.

1.1 Our results

Our main result is a PTAS for clustering problems on graphs of constant highway dimension. For convenience, we define slightly more general problems than those stated above. The k -CLUSTERING ^{q} problem is defined as follows. An instance \mathcal{I} consists of a metric (V, dist) , a set of *facilities* (or *centers*) $F \subseteq V$, and a *demand function* $\chi : V \rightarrow \mathbb{N}_0$. The goal is to find a set $S \subseteq F$ with $|S| \leq k$ minimizing $\sum_{v \in V} \chi(v) \cdot \min_{f \in S} \text{dist}(v, f)^q$. We call all vertices $v \in V$ with $\chi(v) > 0$ the *clients* of \mathcal{I} . k -MEDIAN and k -MEANS are special cases of k -CLUSTERING ^{q} , where $q = 1$ and $q = 2$.

The input to the FACILITY LOCATION ^{q} problem is the same as for k -CLUSTERING ^{q} , but additionally each facility $f \in F$ has an *opening cost* $w_f \in \mathbb{R}^+$. The goal is to find a set $S \subseteq F$ minimizing $\sum_{f \in S} w_f + \sum_{v \in V} \chi(v) \cdot \min_{f \in S} \text{dist}(v, f)^q$. FACILITY LOCATION is a special case of FACILITY LOCATION ^{q} , where $q = 1$.

Our main theorem is the following, where $X = \max_{v \in V} \chi_{\mathcal{I}}(v)$ is the largest demand (note that for k -MEDIAN, k -MEANS, or FACILITY LOCATION we typically have $X = 1$).

► **Theorem 2.** *For any $\varepsilon > 0$, a $(1 + \varepsilon)$ -approximation for k -CLUSTERING ^{q} and FACILITY LOCATION ^{q} can be computed in $(nX)^{(hq/\varepsilon)^{O(q)}}$ time on graphs of highway dimension h .*

In particular, this algorithm is much faster than the quasi-polynomial time approximation scheme of Feldmann et al. [14] for k -MEDIAN or FACILITY LOCATION. The runtime of our algorithm also significantly improves over the exponential dependence on k in the approximation schemes of Becker et al. [6], Braverman et al. [8] for k -MEDIAN.

It has so far been open whether these clustering problems are NP-hard on graphs of constant highway dimension. We complement our main theorem by showing that they are NP-hard even for the smallest possible highway dimension. This answers an open problem given in [14]. Here the *uniform* FACILITY LOCATION ^{q} problem has unit opening costs for all facilities.

► **Theorem 3.** *The k -CLUSTERING^q and uniform FACILITY LOCATION^q problems are NP-hard on graphs of highway dimension 1.*

1.2 Related work

On clustering problems. The problems we focus on in this article are known to be APX-hard, even in Euclidean spaces (see e.g. [4]). In general metric spaces, the current best polynomial-time algorithm for FACILITY LOCATION achieves a 1.488-approximation [19], while the best approximation factor is 2.67 for k -MEDIAN ([9]) and 6.357 for k -MEANS [2].

When restricting the class of graphs, a near-linear time approximation scheme for doubling metrics was developed in [10]; we will discuss the close relations between our work and this one in Section 1.3. Local search techniques also yield a PTAS in minor-free graphs or with bounded doubling dimension [11, 15], and a $\Theta(q)$ -approximation for the k -CLUSTERING^q problem in general metric spaces [17].

Another technique for dealing with clustering problems is to compute *coresets*, a compressed representation of the input. An ε -coreset is a weighted set of points such that for every set of centers, the cost for the original set of points is within a $(1 + \varepsilon)$ -factor of the cost for the coreset. Braverman et al. [8] recently proved that graphs with highway dimension h admit coreset of size $\tilde{O}((k + h)^{O(1/\varepsilon)})$. This enables to compute a $(1 + \varepsilon)$ -approximation by enumerating all possible solutions of the coreset. However, this coreset does not have small highway dimension,¹ and thus cannot be used to boost our algorithms.

On highway dimension. The highway dimension was originally defined by Abraham et al. [1], who specifically chose balls of radius $4r$ in the Definition 1. Since the original definition in [1], several other definitions have been proposed. In particular, Feldmann et al. [14] proved that when choosing a radius cr in Definition 1 for any constant c strictly larger than 4, it is possible to exploit the structure of graphs with constant highway dimension in order to obtain a QPTAS for problems such as TSP, FACILITY LOCATION, and STEINER TREE. As Abraham et al. [1] point out, the choice of the constant is somewhat arbitrary, and we use the above definition so that we may exploit the structural insights of [14] for our algorithm. These structural properties were also leveraged by Becker et al. [6] who gave a PTAS for the BOUNDED-CAPACITY VEHICLE ROUTING problem, and a parameterized approximation scheme for the k -CENTER problem (which is essentially k -CLUSTERING^q with $q = \infty$) and k -MEDIAN. In the lower bound side, Disser et al. [12] showed that STEINER TREE and TSP are weakly NP-hard even when the highway dimension is 1, i.e., each of them is NP-hard but an FPTAS exists for graphs of highway dimension 1.

It is worth mentioning that further definitions of the highway dimension exist (for a detailed discussion see [7, 14]). In particular, for a more general definition of the highway dimension than the one of Definition 1, Feldmann [13] gave a parameterized $3/2$ -approximation algorithm with runtime $2^{O(kh \log h)} n^{O(1)}$ for k -CENTER.

1.3 Our techniques

To obtain Theorem 2, we rely on the framework recently developed by Cohen-Addad et al. [10] for doubling metrics. More precisely, they show that the *split-tree decomposition* of Talwar [20] has some interesting properties, and exploit them to design their algorithm.

¹ Indeed, a subset of a metric with small highway dimension does not necessarily have small highway dimension as well: think of a star metric on which the center is removed.

Our main contribution is to provide a decomposition with similar properties in graphs with constant highway dimension. This is done relying on some structural properties of such graphs presented by Feldmann et al. [14]. We start by giving the outline of the algorithm from [10], and then explain how to carry the results over to the highway dimension setting.

On doubling metrics. The starting point of many approximation algorithms for doubling metrics is a decomposition of the metric, as presented in the following lemma.² A *hierarchical decomposition* \mathcal{D} of a metric (V, dist) is a set of partitions $\mathcal{B}_0, \mathcal{B}_1, \dots, \mathcal{B}_\lambda$, where \mathcal{B}_i refines \mathcal{B}_{i+1} , i.e., every part $B \in \mathcal{B}_i$ is contained in some part of \mathcal{B}_{i+1} . Moreover, in \mathcal{B}_0 every part contains a singleton vertex, while \mathcal{B}_λ contains only one part, namely V . For a point $v \in V$ and a radius $r > 0$, we say that the ball $\beta_v(r)$ is *cut at level i* if i is the largest integer for which the ball $\beta_v(r)$ is not contained in a single part of \mathcal{B}_i . For any subset $W \subseteq V$ we define $\lambda(W) = \lceil \log_2 \text{diam}(W) \rceil$.

► **Lemma 4** (Reformulation of [20, 5] as found in [10]). *For any metric (V, dist) of doubling dimension d and any $\rho > 0$, there exists a polynomial-time computable randomized hierarchical decomposition $\mathcal{D} = \{\mathcal{B}_0, \dots, \mathcal{B}_{\lambda(V)}\}$ such that:*

1. **Scaling probability:** for any $v \in V$, radius r , and level i , we have $\Pr[\mathcal{D} \text{ cuts } \beta_v(r) \text{ at level } i] \leq 2^{O(d)} \cdot r/2^i$.
2. **Portal set:** every part $B \in \mathcal{B}_i$ where $\mathcal{B}_i \in \mathcal{D}$ comes with a set of portals $P_B \subseteq B$ that is
 - a. **concise:** the size of the portal set is bounded by $|P_B| \leq 1/\rho^d$, and
 - b. **precise:** for every node $u \in B$ there is a portal $p \in P_B$ with $\text{dist}(u, p) \leq \rho 2^{i+1}$.

We sketch briefly the standard use of this decomposition. For clustering problems, one can show that there exists a *portal-respecting solution* with near-optimal cost (see Talwar [20]). In this structured solution, each client connects to a facility via a *portal-respecting path* that enters and leaves any part B of \mathcal{D} only through a node of the portal set P_B . Those portals therefore act as separators of the metric. A standard dynamic program approach can then compute the best portal respecting solution.

To ensure that there is a portal-respecting solution with near-optimal cost, one uses the preciseness property of the portal set: the distortion of connecting a client c with a facility f through portals instead of directly is bounded as follows. Let i be the level at which \mathcal{D} cuts c and f , meaning that i is the maximum integer for which c and f lie in different parts of \mathcal{B}_i . At every level $j \leq i$, the distortion incurred by using portals is $\rho 2^j$. Hence the total distortion is $\sum_{j \leq i} \rho 2^j = \rho 2^{i+1}$. Now, property (1) of the decomposition ensures that c and f are cut at level i with probability $O(\text{dist}(c, f)/2^i)$. Hence combining those two bounds over all levels ensures that, in expectation, the distortion between c and f is $O(\text{dist}(c, f) \cdot \rho \lambda(V))$. Since $\lambda(V) = O(\log n)$, choosing $\rho = \varepsilon / \log n$ gives a distortion of $O(\varepsilon \text{dist}(c, f))$. Summing over all clients proves that there exists a near-optimal portal-respecting solution.

The issue with this approach is that the number of needed portals is $O(\log^d n)$, and the dynamic program has a runtime that is exponential in this number. Thus the time complexity is quasipolynomial. The novelty of [10] is to show how to reduce the number of portals to a constant. The idea is to reduce the number of levels on which a client can be cut from its facility.

² We remark that in [10] the preciseness of Lemma 4 was expressed akin to the weaker property found in Lemma 5, which however would not lead to a near-linear time approximation scheme as claimed in [10], but rather a PTAS as shown in this work. This can however easily be alleviated for [10] by using the stronger preciseness as stated here in Lemma 4.

For this, they present a processing step of the instance, that helps deal with clients cut from their facility at a high level. Roughly speaking, their algorithm computes a constant factor approximation L , and a client c is called *badly-cut* if \mathcal{D} cuts it from its closest facility of L at a level larger than $\log(\text{dist}(c, L)/\varepsilon)$. Every badly-cut client is moved to its closest facility of L . Moreover, every client at distance less than $\varepsilon \text{dist}(c, L)$ of its closest facility of L can be moved to it as well. It is then shown that this new instance $\mathcal{I}_{\mathcal{D}}$ has *small distortion*, which essentially means that any solution to $\mathcal{I}_{\mathcal{D}}$ can be converted to a solution of the original instance \mathcal{I} while only losing a $(1 + \varepsilon)$ -factor in quality. In this instance $\mathcal{I}_{\mathcal{D}}$, all clients are cut from their closest facility of L at some level between $\log(\varepsilon \text{dist}(c, L))$ and $\log(\text{dist}(c, L)/\varepsilon)$. Using this property, it can be shown that c and its closest center in the optimal solution are also cut at a level in that range. As there are only $O(\log(1/\varepsilon))$ levels in this range, by the previous argument, the number of portals is a constant. (See Section 2 for formal definitions and lemmas.)

On highway dimension. The above arguments for doubling metrics hold thanks to Lemma 4. In this work, we show how to construct a similar decomposition for low highway dimension:

► **Lemma 5.** *Given a shortest-path metric (V, dist) of a graph with highway dimension h , a subset $W \subseteq V$, and $\rho > 0$, there exists a polynomial-time computable randomized hierarchical decomposition $\mathcal{D} = \{\mathcal{B}_0, \dots, \mathcal{B}_{\lambda(W)}\}$ of W such that:*

1. **Scaling probability:** for any $v \in V$, radius r , and level i , we have $\Pr[\mathcal{D} \text{ cuts } \beta_v(r) \text{ at level } i] \leq \sigma \cdot r/2^i$, where $\sigma = (h \log(1/\rho))^{O(1)}$.
2. **Interface:** for any $B \in \mathcal{B}_i$ on level $i \geq 1$ there exists an interface $I_B \subseteq V$, which is
 - a. **concise:** $|I_B| \leq (h/\rho)^{O(1)}$, and
 - b. **precise:** for any $u, v \in B$ such that u and v are cut by \mathcal{D} at level $i - 1$, there exists $p \in I_B$ with $\text{dist}(u, p) + \text{dist}(p, v) \leq \text{dist}(u, v) + 34 \cdot \rho 2^i$.

Our construction relies on the *town decomposition* from [14], which has the following properties: for a graph with highway dimension h and a given $\rho > 0$, every part T of the decomposition (called a *town*) has a set X_T of *hubs* with doubling dimension $O(h \log 1/\rho)$, such that for any two vertices u and v in different child towns of T , there is a vertex $x \in X_T$ such that $\text{dist}(u, x) + \text{dist}(x, v) \leq (1 + 2\rho) \cdot \text{dist}(u, v)$ – see Theorem 8 for more details.

This hub set X_T is similar to the portal set of Lemma 4, but has some fundamental differences: the first one is that the decomposition is deterministic, and so it may happen that a client and its facility are cut at a very high level – something that happens only with tiny probability in the doubling setting thanks to the scaling probability. Another main difference is that the size of X_T might be unbounded. As a consequence, it cannot be directly used as a portal set in a dynamic program. To deal with this, we combine the town decomposition with a hierarchical decomposition of each set X_T according to Lemma 4, to build an *interface* as stated in Lemma 5.

A further notable difference to portals is that the preciseness property of the resulting interface is weaker. In particular, while there is a portal close to each vertex of a part, the hubs can be far from some vertices as long as they lie close to the shortest path to other vertices, which however can be far (due to Lemma 9). As a consequence no analog of near-optimal portal-respecting paths exist. Instead, when connecting a client c with a facility f we need to use the interface point of I_B provided by the preciseness property of Lemma 5 close to the shortest path between c and f , where B contains both c and f . This shifts the perspective from externally connecting vertices of a part to vertices outside a part, as done for portals, to internally connecting vertices of parts, as done here.

As a consequence, we develop a dynamic program, which follows more or less standard techniques as for instance given in [3, 18], but needs to handle the weaker preciseness property of the interface. The main idea is to guess the distances from interface points to facilities while recursing on the decomposition \mathcal{D} of Lemma 5. Due to the shifted perspective towards internally connecting vertices of parts, the runtime of the dynamic program depends exponentially on the *total* number of levels. However, it can be shown that it suffices to compute a solution on a carefully chosen subset W of the metric for which only a logarithmic number of levels of the decomposition need to be considered, and thus the runtime is polynomial.

1.4 Outline

After defining the concepts we use, and stating various structural lemmas in Section 2, we show how to incorporate our decomposition into the framework of [10]. The proof of Lemma 5 is then presented in Section 3. The formal algorithm is deferred to Section 4. We conclude the main body of this paper with the hardness proof of Theorem 3 in Section 5.

2 Preliminaries

On doubling metrics. The *doubling dimension* of a metric is the smallest integer d such that for any $r > 0$ and $v \in V$, the ball $\beta_v(2r)$ of radius $2r$ around v can be covered by at most 2^d balls of half the radius r . A doubling metric is a metric space where the doubling dimension is bounded. In those spaces, one can show the existence of small *nets*:

► **Definition 6.** A δ -net of a metric (V, dist) is a subset of nodes $N \subseteq V$ with the property that every node in V is at distance at most δ from a net point of N , and each pair of net points of N are at distance more than δ .

► **Lemma 7** ([16]). Let (V, dist) be a metric space with doubling dimension d . If its diameter is D , and N is a δ -net of V , then $|N| \leq 2^{d \cdot \lceil \log_2(D/\delta) \rceil}$. Moreover, any subset $W \subseteq V$ has doubling dimension at most $2d$.

On highway dimension. We note that for simplicity we will set $c = 8$ in Definition 1 throughout this paper, even if all claimed results are also true for other values of c . When we refer to a metric as having highway dimension h , we mean that it is the shortest-path metric of a graph of highway dimension h . The main result we will use about highway dimension is existence the of the following decomposition:

► **Theorem 8** ([14]). Given a shortest-path metric (V, dist) of highway dimension h , and $\rho > 0$, there exists a polynomial-time computable deterministic hierarchical decomposition \mathcal{T} , called the town decomposition, such that every part $T \in \mathcal{T}$, called a town, has a set of hubs³ $X_T \subseteq T$ with the following properties:

- a. **doubling:** the doubling dimension of X_T is $d = O(\log(h \log(1/\rho)))$, and
- b. **precise:** for any two vertices u and v in different child parts of T , there is a vertex $x \in X_T$ such that $\text{dist}(u, x) + \text{dist}(x, v) \leq (1 + 2\rho) \cdot \text{dist}(u, v)$.

The town decomposition behaves differently from those in Lemmas 4 and 5 in several ways. The main properties we will need here are the following.

► **Lemma 9** ([14]). For any $T \in \mathcal{T}$ we have $\text{diam}(T) < \text{dist}(T, V \setminus T)$. Furthermore, for any child town T' of T we have $\text{diam}(T') \leq \text{diam}(T)/2$.

³ called *approximate core hubs* in [14].

On how to incorporate our decomposition into the framework of [10]. Assume we are given an instance \mathcal{I} of k -CLUSTERING^q or FACILITY LOCATION^q on some metric (V, dist) , together with a hierarchical decomposition \mathcal{D} of the metric with the properties listed in Lemma 5. We start by defining the *badly cut* clients. In the following, we fix an optimal solution OPT and an approximate solution L , and we define $\tau(\varepsilon, q, \sigma) = \log_2(\sigma(q+1)^q/\varepsilon^{q+1})$.

► **Definition 10** (badly cut [10]). *Let (V, dist) be a metric of an instance \mathcal{I} of k -CLUSTERING^q or FACILITY LOCATION^q, \mathcal{D} be a hierarchical decomposition of the metric with scaling probability factor σ , and $\varepsilon > 0$. If L_v is the distance from v to the closest facility of an approximate solution L to \mathcal{I} , then a client c is badly cut w.r.t. \mathcal{D} if the ball $\beta_c(3L_c/\varepsilon)$ is cut as some level i greater than $\log_2(3L_c/\varepsilon) + \tau(\varepsilon, q, \sigma)$.*

Similarly, if OPT_v is the distance from v to the closest facility of the optimum solution OPT of \mathcal{I} , then a facility $f \in L$ is badly cut w.r.t. \mathcal{D} if $\beta_f(3\text{OPT}_f)$ is cut at some level i greater than $\log_2(3\text{OPT}_f) + \tau(\varepsilon, q, \sigma)$.

Given an instance \mathcal{I} of k -CLUSTERING^q or FACILITY LOCATION^q and a decomposition \mathcal{D} of the metric, a new instance $\mathcal{I}_{\mathcal{D}}$ is computed to get rid of badly cut clients. The instance $\mathcal{I}_{\mathcal{D}}$ is built from \mathcal{I} by moving clients that are badly cut w.r.t. \mathcal{D} to their closest facility in L .⁴ For any client c of $\mathcal{I}_{\mathcal{D}}$ we denote by \tilde{c} the original position of this client in \mathcal{I} , i.e., if \tilde{c} is a badly cut client of \mathcal{I} then $c = L(\tilde{c})$ and otherwise $c = \tilde{c}$. The set F of potential centers is unchanged, and thus any solution of \mathcal{I} is a solution of $\mathcal{I}_{\mathcal{D}}$, and vice versa. Note that $\mathcal{I}_{\mathcal{D}}$ does not contain any badly cut client w.r.t. \mathcal{D} , and that the definition of $\mathcal{I}_{\mathcal{D}}$ depends on the randomness of \mathcal{D} .

To describe the properties we obtain for the new instance, given a solution S to any instance \mathcal{I}_0 of k -CLUSTERING^q or FACILITY LOCATION^q, we define $\text{cost}_{\mathcal{I}_0}(S) = \sum_{v \in V} \chi_{\mathcal{I}_0}(v) \cdot \text{dist}(v, S)^q$ to be the cost incurred by only the distances to the facilities. Given some $\varepsilon > 0$ and the computed instance $\mathcal{I}_{\mathcal{D}}$ from \mathcal{I} , we define

$$\nu_{\mathcal{I}_{\mathcal{D}}} = \max_{\text{solution } S} \left\{ \text{cost}_{\mathcal{I}}(S) - (1 + 2\varepsilon) \text{cost}_{\mathcal{I}_{\mathcal{D}}}(S), (1 - 2\varepsilon) \text{cost}_{\mathcal{I}_{\mathcal{D}}}(S) - \text{cost}_{\mathcal{I}}(S) \right\}.$$

If $B_{\mathcal{D}}$ denotes the set of badly cut facilities (w.r.t. \mathcal{D}) of the solution L to \mathcal{I} from which instance $\mathcal{I}_{\mathcal{D}}$ is constructed, we say that $\mathcal{I}_{\mathcal{D}}$ has *small distortion w.r.t. \mathcal{I}* if $\nu_{\mathcal{I}_{\mathcal{D}}} \leq \varepsilon \text{cost}_{\mathcal{I}}(L)$, and there exists a *witness solution* $\hat{S} \subseteq F$ that contains $B_{\mathcal{D}}$ and for which $\text{cost}_{\mathcal{I}_{\mathcal{D}}}(\hat{S}) \leq (1 + O(\varepsilon)) \text{cost}_{\mathcal{I}}(\text{OPT}) + O(\varepsilon) \text{cost}_{\mathcal{I}}(L)$. Moreover, in the case of FACILITY LOCATION^q, $\hat{S} = \text{OPT} \cup B_{\mathcal{D}}$ and $\sum_{f \in B_{\mathcal{D}}} w_f \leq \varepsilon \cdot \sum_{f \in L} w_f$.

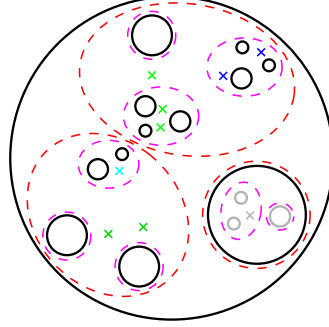
Based on these definitions, we now state the main tool we use from [10], and which exploits the scaling probability of our decomposition in Lemma 5 to obtain the required structure.

► **Lemma 11** ([10]). *Let (V, dist) be a metric, and \mathcal{D} be a randomized hierarchical decomposition of (V, dist) with scaling probability factor σ . Let \mathcal{I} be an instance of k -CLUSTERING^q or FACILITY LOCATION^q on (V, dist) , with optimum solution OPT and approximate solution L . For any (sufficiently small) $\varepsilon > 0$, with probability at least $1 - \varepsilon$ (over \mathcal{D}), the instance $\mathcal{I}_{\mathcal{D}}$ constructed from \mathcal{I} and L as described above has small distortion with a witness solution \hat{S} . Furthermore, every client c of $\mathcal{I}_{\mathcal{D}}$ is cut by \mathcal{D} from its closest facility in \hat{S} at level at most $\log_2(3L_{\tilde{c}}/\varepsilon + 4\text{OPT}_{\tilde{c}}) + \tau(\varepsilon, q, \sigma)$, where \tilde{c} is the original position of c in \mathcal{I} .*

As a consequence of Lemma 11, a dynamic program can compute a solution recursively on the parts of \mathcal{D} in polynomial time, as sketched in Section 1.3 and detailed in Section 4.

⁴ More concretely, let $\chi_{\mathcal{I}}$ and $\chi_{\mathcal{I}_{\mathcal{D}}}$ be the demand functions of \mathcal{I} and $\mathcal{I}_{\mathcal{D}}$, respectively. Initially we let $\mathcal{I}_{\mathcal{D}}$ be a copy of \mathcal{I} , so that in particular $\chi_{\mathcal{I}_{\mathcal{D}}} = \chi_{\mathcal{I}}$. Then, for each client c of \mathcal{I} that is badly cut in L w.r.t. \mathcal{D} , if $L(c)$ denotes the closest facility of L to c , in $\mathcal{I}_{\mathcal{D}}$ we set $\chi_{\mathcal{I}_{\mathcal{D}}}(c) = 0$ and increase $\chi_{\mathcal{I}_{\mathcal{D}}}(L(c))$ by the value of $\chi_{\mathcal{I}}(c)$ in \mathcal{I} .

3 Decomposing the graph



■ **Figure 1** A town T and its child towns (black circles). The hubs (crosses) are decomposed by \mathcal{X}_T (indicated by different colours). Parts $B \in \mathcal{B}_{i+1}$ (red dashed) are decomposed into parts on level i (pink dashed). Parts of \mathcal{B}_{i-1} can lie in different towns (e.g., the child town of T with subtowns in grey).

This section is dedicated to the proof of Lemma 5. The general idea to construct \mathcal{D} is as follows. For doubling metrics, to decompose a part at level i , it is enough to pick a random diameter $\delta \in [2^{i-2}, 2^{i-1})$ and divide the part into child parts of diameter δ . This is not doable in the highway dimension setting: if one wishes to decompose a town T , it cannot divide any of the child towns, since it is not possible to use the approximate core hubs of T to approximate paths inside one of the child towns. The big picture of our decomposition is therefore as follow. To decompose a town at level i , we group randomly (as in the doubling decomposition) the “small” child towns, and put every “big” child town in its own subpart. As we will see, this turns out to be enough.

In order to decompose a town T , we need the following definitions. For each child town T' of T we identify the *connecting hub* $x \in X_T$, which is some fixed closest hub of X_T to T' , breaking ties arbitrarily. Moreover, given a hierarchical decomposition $\mathcal{X}_T = \{\mathcal{U}_0, \dots, \mathcal{U}_{\lambda(X_T)}\}$ of X_T , we define for every i the *connecting i -cluster* of a child town T' of T to be the set $U \in \mathcal{U}_\ell$ on level $\ell = \min\{i, \lambda(X_T)\}$ containing the connecting hub of T' . We then follow the steps below, after choosing μ from the interval $(0, 1]$ uniformly at random (cf. Figure 1):

1. Given a town $T \in \mathcal{T}$, we apply Lemma 4 to find a randomized hierarchical decomposition $\mathcal{X}_T = \{\mathcal{U}_0, \dots, \mathcal{U}_{\lambda(X_T)}\}$ of the hubs X_T of T .
2. Using \mathcal{X}_T , we define a randomized partial decomposition of $T \cap W$ as follows. For any i and $U \in \mathcal{U}_{\min\{i, \lambda(X_T)\}}$, let the set $A_i^U \subseteq T \cap W$ be the union of all $T' \cap W$ where T' is a child town of T with the following two properties:
 - a. U is the connecting i -cluster of T' , and
 - b. $\text{dist}(T', V \setminus T') \leq \mu 2^i$.
Hence A_i^U contains all towns somewhat close to U , and with small diameter due to Lemma 9. We let \mathcal{A}_i^T be the set containing each non-empty A_i^U .
3. Now, the hierarchical decomposition $\mathcal{D} = \{\mathcal{B}_0, \dots, \mathcal{B}_{\lambda(W)}\}$ of W can be constructed inductively as follows. At the highest level $\lambda(W)$ of \mathcal{D} , W is partitioned in a single set: $\mathcal{B}_{\lambda(W)} = \{W\}$. Now, to decompose a part $B \in \mathcal{B}_{i+1}$ at level $i+1$, we do the following. Let $T \in \mathcal{T}$ be the inclusion-wise minimal town for which $B \subseteq T$. The “small” subtowns of T lying inside B are grouped according to step (2) (note that $\text{dist}(T', V \setminus T')$ also bounds the diameter of T' by Lemma 9), and the other ones form individual subparts. More formally, the set \mathcal{B}_i contains every part $A \in \mathcal{A}_i^T$ for which $A \subseteq B$, and also every set $T' \cap W$, where T' is a child town of T for which $T' \cap W \subseteq B$ and $T' \cap W$ was not covered by the previously added parts of \mathcal{A}_i^T , i.e., $T' \cap W \cap A = \emptyset$ for every $A \in \mathcal{A}_i^T$.

To prove that the constructed decomposition \mathcal{D} has the desired properties –i.e. that it is indeed a hierarchical decomposition, with parts of bounded diameter and small scaling probability factor –, we begin with some auxiliary lemmas, of which the first one bounds the distance of a town to its connecting hub.

► **Lemma 12.** *If T' is a child town of T with connecting hub $x \in X_T$, then $\text{dist}(x, T') \leq (1 + 2\rho) \text{dist}(T', V \setminus T')$.*

Proof. Let T'' be the closest sibling town to T' , and let $u \in T'$ and $v \in T''$ be the vertices defining the distance from T' to T'' , i.e., $\text{dist}(u, v) = \text{dist}(T', T'') = \text{dist}(T', V \setminus T')$. By Theorem 8, there is a hub $y \in X_T$ for which $\text{dist}(u, y) + \text{dist}(y, v) \leq (1 + 2\rho) \cdot \text{dist}(u, v) = (1 + 2\rho) \cdot \text{dist}(T', V \setminus T')$. This implies $\text{dist}(y, T') \leq \text{dist}(u, y) \leq (1 + 2\rho) \cdot \text{dist}(T', V \setminus T')$. Since the connecting hub x of T' is at least as close to T' as y , the claim follows. ◀

Based on the above lemma, we next prove a key property that the diameter of any part of $\mathcal{B}_i \in \mathcal{D}$ is bounded.

► **Lemma 13.** *If $\rho \leq 1/2$, then the diameter of any part of $\mathcal{B}_i \in \mathcal{D}$ is less than 2^{i+4} .*

Proof. On the highest level $\lambda(W)$ of \mathcal{D} the only part of $\mathcal{B}_{\lambda(W)}$ is W itself. As $\lambda(W) = \lceil \log_2 \text{diam}(W) \rceil$ we get $\text{diam}(W) \leq 2^{\lambda(W)+1}$, as required.

For any level $i < \lambda(W)$, a set in \mathcal{B}_i is equal to a set $A \in \mathcal{A}_i^T$ for some town $T \in \mathcal{T}$ or it is equal to some set $T' \cap W$ for a child town T' of T . In the former case, the set A is equal to a set A_i^U for some cluster $U \in \mathcal{U}_\ell$ where $\ell = \min\{i, \lambda(X_T)\}$ and $\mathcal{U}_\ell \in \mathcal{X}_T$. The set A_i^U contains the union of sets $T' \cap W$ for child towns T' of T , for which their connecting hubs lie in U and $\text{dist}(T', V \setminus T') \leq \mu 2^i \leq 2^i$, as $\mu \leq 1$. Thus from Lemma 12 we get $\text{dist}(U, T') \leq (1 + 2\rho) 2^i$, and by Lemma 9 we have $\text{diam}(T') < \text{dist}(T', V \setminus T') \leq 2^i$. The cluster U has diameter less than 2^{i+1} by Lemma 4, since it is part of the hierarchical decomposition \mathcal{X}_T and lies on level $\ell \leq i$. Let u and v be the vertices of A_i^U defining the diameter of A_i^U , i.e., $\text{dist}(u, v) = \text{diam}(A_i^U)$. We may reach v from u by first crossing the child town T' that u lies in, then passing over to U , then crossing U , after which we pass over to the child town T'' containing v , and finally crossing this child town as well to reach v . Hence, assuming that $\rho \leq 1/2$ the diameter of A_i^U is bounded by

$$\begin{aligned} \text{dist}(u, v) &\leq \text{diam}(T') + \text{dist}(U, T') + \text{diam}(U) + \text{dist}(U, T'') + \text{diam}(T'') \\ &< 2 \cdot 2^i + 2 \cdot (1 + 2\rho) 2^i + 2^{i+1} = (6 + 4\rho) 2^i \leq 2^{i+3} \end{aligned}$$

Now consider the other case, when a set $B \in \mathcal{B}_i$ on level $i < \lambda(W)$ is equal to some set $T' \cap W$ for a child town T' of a town T . For such a child town T' there is no enforced upper bound on the distance to other child towns as before, and thus it is necessary to be more careful to bound the diameter of the part. Starting with $B = B_i$, let $B_i \subseteq B_{i+1} \subseteq \dots \subseteq B_j$ be the longest chain of parts of increasing levels that are of the same type as B . More concretely, for every $\ell \in \{i, i+1, \dots, j\}$ we have $B_\ell \in \mathcal{B}_\ell$ and B_ℓ is equal to some set $T'_\ell \cap W$ for a child town T'_ℓ of the inclusion-wise minimal town T_ℓ containing $B_{\ell+1}$. Note that in particular $j < \lambda(W)$. As we chose the longest such chain, on the next level $j+1$ there is no such set containing B_j , which means that the set $B_{j+1} \in \mathcal{B}_{j+1}$ for which $B_j \subseteq B_{j+1}$ is either equal to a set $A \in \mathcal{A}_{j+1}^{T_{j+1}}$ for some town T_{j+1} , or $j+1 = \lambda(W)$. In either case, from above we get $\text{diam}(B_{j+1}) \leq 2^{j+4}$.

Note that for any $\ell \in \{i, i+1, \dots, j-1\}$ the town T'_ℓ is a descendant town of $T'_{\ell+1}$, since $B_{\ell+1}$ is contained in $T'_{\ell+1}$ and T'_ℓ is a child town of the inclusion-wise minimal town T_ℓ containing $B_{\ell+1}$. By Theorem 8 and Lemma 9 we thus get $\text{diam}(T'_\ell) \leq \text{diam}(T'_{\ell+1})/2$,

which implies $\text{diam}(T'_i) \leq \text{diam}(T'_j)/2^{j-i}$. The set $B = B_i$ is contained in T'_i , which means $\text{diam}(B) \leq \text{diam}(T'_i)$. The town T_j is the inclusion-wise minimal town containing B_{j+1} , while at the same time the child town T'_j of T_j contains B_j . As $B_j \subseteq B_{j+1}$, this means that B_{j+1} both contains vertices inside and outside of T'_j , and so $\text{dist}(T'_j, V \setminus T'_j) \leq \text{diam}(B_{j+1})$. By Lemma 9 we know that $\text{diam}(T'_j) \leq \text{dist}(T'_j, V \setminus T'_j)$, and putting all these inequalities together we obtain

$$\begin{aligned} \text{diam}(B) &\leq \text{diam}(T'_i) \leq \text{diam}(T'_j)/2^{j-i} \leq \text{dist}(T'_j, V \setminus T'_j)/2^{j-i} \\ &\leq \text{diam}(B_{j+1})/2^{j-i} \leq 2^{j+4}/2^{j-i} = 2^{i+4}. \quad \blacktriangleleft \end{aligned}$$

Using Lemma 13 it is not hard to prove the correctness of \mathcal{D} , which we turn to next. All statements marked with “ \star ” are deferred to the full version of the paper, due to space constraints.

► **Lemma 14** (\star). *The tuple $\mathcal{D} = \{\mathcal{B}_0, \dots, \mathcal{B}_{\lambda(V)}\}$ is a hierarchical decomposition of W .*

We now turn to proving the properties of Lemma 5, starting with the scaling probability.

► **Lemma 15.** *The decomposition \mathcal{D} has scaling probability factor $\sigma = (h \log(1/\rho))^{O(1)}$.*

Proof. To prove the claim, we need to prove that for any $v \in W$, radius r , and level i , the probability that \mathcal{D} cuts the ball $\beta_v(r)$ at level i is at most $(h \log(1/\rho))^{O(1)} \cdot r/2^i$. If \mathcal{D} cuts $\beta_v(r)$ at level i , it means that $\beta_v(r)$ is fully contained in a part at level $i+1$: let $T \in \mathcal{T}$ be the inclusion-wise minimal town containing that part. There are two cases to consider: either $\beta_v(r)$ is cut by “small” parts, i.e. there exist two distinct parts $A, A' \in \mathcal{A}_i^T$ such that $v \in A$ and $u \in A'$ for some $u \in W \cap \beta_v(r)$, or not.

We start with the latter case, when $\beta_v(r)$ is not cut by small parts. If \mathcal{D} cuts the ball at level i , there are distinct parts $B, B' \in \mathcal{B}_i$ such that $v \in B$ and $u \in B'$ for some $u \in W \cap \beta_v(r)$. Assume w.l.o.g. that $B \notin \mathcal{A}_i^T$ (which is possible to assume since $\beta_v(r)$ is not cut by small parts). By construction of the decomposition, there must be a child town T' of T , for which $B = T' \cap W$ and $\text{dist}(T', V \setminus T') > \mu 2^i$. Note that $r \geq \text{dist}(v, u) \geq \text{dist}(T', B'') \geq \text{dist}(T', V \setminus T') \geq \mu 2^i$, and hence $\mu \leq r/2^i$. The decomposition \mathcal{D} can therefore only cut $\beta_v(r)$ on level i if $\mu < r/2^i$. Since μ is chosen uniformly at random from the interval $(0, 1]$, the probability is less than $r/2^i$.

We now turn to the other case, when $\beta_v(r)$ is cut by two small parts A_1 and A_2 . The town T must have two child towns T_1 and T_2 for which $v \in T_1 \cap W \subseteq A_1$ and $u \in T_2 \cap W \subseteq A_2$. Let x_1 and x_2 be the connecting hubs of T_1 and T_2 . The decomposition \mathcal{D} cuts v and u on level i if and only if \mathcal{X}_T cuts x_1 and x_2 on level $\ell = \min\{i, \lambda(X_T)\}$. Indeed, let U_1 and U_2 be the connecting i -clusters of T_1 and T_2 : then $A_1 = A_{U_1}^i$ and $A_2 = A_{U_2}^i$, with $x_1 \in U_1, x_2 \in U_2$. Thus \mathcal{D} cuts v and u on level i if and only if $U_1 \neq U_2$, i.e., if and only if \mathcal{X}_T cuts x_1 and x_2 on level $\ell = \min\{i, \lambda(X_T)\}$.

To compute the probability that x_1 and x_2 are cut, it is necessary to bound the distance between them. As $v \in T_1$ and $u \in T_2$ while $u \in \beta_v(r)$, for each $j \in \{1, 2\}$ we have $\text{dist}(T_j, V \setminus T_j) \leq \text{dist}(T_1, T_2) \leq r$. By Lemma 12 the distance between T_j and its connecting hub $x_j \in X_T$ is thus at most $(1 + 2\rho)r$. Also, by Lemma 9 we have $\text{diam}(T_j) < \text{dist}(T_j, V \setminus T_j) \leq r$, and we get

$$\text{dist}(x_1, x_2) \leq \text{dist}(x_1, T_1) + \text{diam}(T_1) + \text{dist}(T_1, T_2) + \text{diam}(T_2) + \text{dist}(T_2, x_2) \leq (5 + 4\rho)r.$$

We can reformulate the above as follows: if \mathcal{D} cuts the ball $\beta_v(r)$ at level i , and $\beta_v(r)$ is cut by some “small” parts A_1 and A_2 , then \mathcal{X}_T cuts the ball $\beta_{x_1}((5 + 4\rho)r)$ on level i , where x_1 is the hub defined for v above. We know that the probability of the latter event

is at most $2^{O(d)}(5 + 4\rho)r/2^i$ by Lemma 4, where $d = O(\log(h \log(1/\rho)))$ is the doubling dimension of X_T by Theorem 8. Hence the probability that \mathcal{D} cuts the ball $\beta_v(r)$ at level i is bounded by $(h \log(1/\rho))^{O(1)} \cdot r/2^i$. Taking a union bound over the two considered cases proves the claim. \blacktriangleleft

To prove the remaining property of Lemma 5 for \mathcal{D} , for each $B \in \mathcal{B}_i$ we need to choose an interface I_B from the whole vertex set V . For this we use a carefully chosen net (see Definition 6) of the hubs of the inclusion-wise minimal town T containing B , as formalized in the following lemma.

► Lemma 16. *Given $B \in \mathcal{B}_i$ for some $\mathcal{B}_i \in \mathcal{D}$ and $i \geq 1$, let $T \in \mathcal{T}$ be the inclusion-wise minimal town containing B . We define the interface I_B to be a $\rho 2^i$ -net of the set $Y_B = \{x \in X_T \mid \text{dist}(x, B) \leq (1 + 2\rho) \text{diam}(B)\}$. The interface I_B has the conciseness and preciseness properties of Lemma 5 for $\rho \leq 1/2$.*

Proof. We first prove that I_B is precise. Consider two vertices $u, v \in B$ that are cut at level $i - 1$ by \mathcal{D} . This means there are two distinct parts $B', B'' \in \mathcal{B}_{i-1}$ on this level such that $v \in B'$ and $u \in B''$. By definition, both B' and B'' are unions of sets $T' \cap W$ where T' is a child town of the inclusion-wise minimal town T containing B . Also $B' \cap B'' = \emptyset$ by Lemma 14. This means that T has two child towns T_1 and T_2 for which $v \in T_1 \cap W \subseteq B'$ and $u \in T_2 \cap W \subseteq B''$. By Theorem 8, there is an approximate core hub $x \in X_T$ such that $\text{dist}(u, x) + \text{dist}(x, v) \leq (1 + 2\rho) \text{dist}(u, v)$. In particular, $\text{dist}(x, B) \leq \text{dist}(u, x) \leq (1 + 2\rho) \text{dist}(u, v) \leq (1 + 2\rho) \text{diam}(B)$, as $u, v \in B$. This means that $x \in Y_B$. Since I_B is a $\rho 2^i$ -net of Y_B , there is a hub $p \in I_B$ for which $\text{dist}(x, p) \leq \rho 2^i$. By Lemma 13 we have $\text{dist}(u, v) \leq \text{diam}(B) \leq 2^{i+4}$ if $\rho \leq 1/2$, and so I_B is precise:

$$\begin{aligned} \text{dist}(u, p) + \text{dist}(p, v) &\leq \text{dist}(u, x) + 2 \cdot \text{dist}(x, p) + \text{dist}(x, v) \\ &\leq (1 + 2\rho) \text{dist}(u, v) + \rho 2^{i+1} \leq \text{dist}(u, v) + 2\rho \cdot 2^{i+4} + \rho 2^{i+1} \leq \text{dist}(u, v) + 34 \cdot \rho 2^i, \end{aligned}$$

To prove conciseness, recall that $\text{diam}(B) \leq 2^{i+4}$ by Lemma 13, which means that $\text{diam}(Y_B) \leq \text{diam}(B) + 2(1 + 2\rho) \text{diam}(B) \leq 5 \cdot 2^{i+4}$ for $\rho \leq 1/2$. Since I_B is a $\rho 2^i$ -net of Y_B , Lemma 7 implies $|I_B| \leq 2^{d \cdot \lceil \log_2(80/\rho) \rceil}$, where d is the doubling dimension of Y_B . Theorem 8 says that X_T has doubling dimension $O(\log(h \log(1/\rho)))$, and as $Y_B \subseteq X_T$ the same asymptotic bound holds for the doubling dimension d of Y_B by Lemma 7. Therefore we get $|I_B| \leq (h \log(1/\rho))^{O(\log(1/\rho))} = (h/\rho)^{O(1)}$, which concludes the proof. \blacktriangleleft

4 The algorithm

Let an instance \mathcal{I} of the k -CLUSTERING^q or FACILITY LOCATION^q problem on a shortest-path metric (V, dist) of a graph G with highway dimension h , and maximum demand $X = \max_{v \in V} \chi_{\mathcal{I}}(v)$ be given. The algorithm performs the following steps:

1. compute a town decomposition \mathcal{T} together with the hubs for each town as given by Theorem 8.
2. compute a hierarchical decomposition \mathcal{D} according to Lemma 5, while simultaneously converting \mathcal{I} into a *coarse* instance w.r.t. \mathcal{D} , meaning that there is a subset $W \subseteq V$ for which
 - the clients and facilities of \mathcal{I} are contained in W , i.e., $F \cup \{v \in V \mid \chi_{\mathcal{I}}(v) > 0\} \subseteq W$, and
 - every part of \mathcal{D} on level at most $\xi(W) = \lfloor \lambda(W) - 2 \log_2(nX/\varepsilon) \rfloor$ has at most one facility, i.e., $|B \cap F| \leq 1$ for every $B \in \mathcal{B}_{\xi(W)}$.

3. compute the instance $\mathcal{I}_{\mathcal{D}}$ of small distortion as given by Lemma 11.
4. run a dynamic program on $\mathcal{I}_{\mathcal{D}}$ as given in Section 4.2, to compute an *optimum rounded interface-respecting solution* (see Section 4.1 for a formal definition), and output it as a solution to the input instance.

In a nutshell, the coarseness of the instances guarantees that only a logarithmic number of levels need to be considered by the dynamic program. This step loses a $(1 + \varepsilon)$ -factor in the solution quality. The dynamic program is only able to compute highly structured solutions, which are captured by the notion of *rounding* and *interface-respecting*. Due to this, another $(1 + \varepsilon)$ -factor in the solution quality is lost. In Section 4.1 we prove that the output of the dynamic program is a near-optimal solution to the input instance (proving Theorem 2), and we also detail step (4) of the algorithm. Then in Section 4.2 we describe the details of the dynamic program.

4.1 Approximating the distances

One caveat of the dynamic program is that the runtime is only polynomial if the recursion depth is logarithmic. However when computing our decomposition on the whole metric (V, dist) , the number of levels is $\lambda(V) + 1 = \lceil \log_2 \text{diam}(V) \rceil + 1$, which can be linear in the input size. Standard techniques can be used to reduce the number of levels to $O(\log(n/\varepsilon))$ when aiming for a $(1 + \varepsilon)$ -approximation by preprocessing the input metric. However, for graphs of bounded highway dimension these general techniques change the hub sets and we would have to be careful to maintain the properties we need, as given by Theorem 8. Therefore we adapt the standard techniques to our setting via the notion of *coarse* instances.

The following lemma shows that we can reduce any instance to a set of coarse ones, for which, as we will see, our dynamic program only needs to consider the highest $2 \log_2(nX/\varepsilon)$ levels.

► **Lemma 17** (\star). *Let \mathcal{I} be an instance of k -CLUSTERING^q or FACILITY LOCATION^q on a graph G of highway dimension h . There are polynomial-time computable instances $\mathcal{I}_1, \dots, \mathcal{I}_b$ and respective hierarchical decompositions $\mathcal{D}_1, \dots, \mathcal{D}_b$ with the properties given in Lemma 5 for any $\rho \leq 1/2$, such that for each $i \in \{1, \dots, b\}$ the instance \mathcal{I}_i is also defined on G and is coarse w.r.t. \mathcal{D}_i . Furthermore, if an α -approximation can be computed for each of the instances $\mathcal{I}_1, \dots, \mathcal{I}_b$ in polynomial time, then for any $\varepsilon > 0$ a $(1 + O(\varepsilon))\alpha$ -approximation can be computed for \mathcal{I} in polynomial time.*

Lemma 17 implies that if there is a PTAS for coarse instances, we also have a PTAS in general. Hence from now on we assume that the given instance \mathcal{I} is coarse w.r.t. a hierarchical decomposition \mathcal{D} of some subset W of the vertices of the input graph G , where \mathcal{D} has bounded scaling probability factor and concise and precise interface sets in G according to Lemma 5 (for some value $\rho > 0$ specified later)

The next step of the algorithm is to compute a new instance $\mathcal{I}_{\mathcal{D}}$ with small distortion as given by Lemma 11. Recall that $\mathcal{I}_{\mathcal{D}}$ is obtained from \mathcal{I} by moving badly cut clients to facilities of L . In particular, the instance $\mathcal{I}_{\mathcal{D}}$ is also coarse w.r.t. \mathcal{D} , which means that we may run our dynamic program on $\mathcal{I}_{\mathcal{D}}$.

The dynamic program exploits the interface sets of \mathcal{D} by computing a near-optimum “interface-respecting” solution to $\mathcal{I}_{\mathcal{D}}$, i.e., a solution where clients are connected to facilities through interface points. Moreover, for the dynamic program to run in polynomial time it can only estimate the distances between interface points and facilities to a certain precision. In general, we denote by $\langle x \rangle_i = \min\{(35 + \delta)\rho 2^i \mid \delta \in \mathbb{N} \text{ and } \rho \delta 2^i \geq x\}$ the value of x rounded to

the next multiple of $\rho 2^i$ and shifted by $35\rho 2^i$. We then define the *rounded interface-respecting distance* $\text{dist}'(v, u)$ from a vertex v to another vertex u as follows. If v and u are not cut at any level, i.e., $v = u$, then $\text{dist}'(v, u) = 0$. Otherwise, if $i \geq 1$ is the level of \mathcal{D} such that there is a part $B \in \mathcal{B}_i$ with $v, u \in B$, and \mathcal{D} cuts v and u at level $i - 1$, we let

$$\text{dist}'(v, u) = \min \{ \text{dist}(v, p) + \langle \text{dist}(p, u) \rangle_i \mid p \in I_B \}.$$

Note that $\text{dist}'(\cdot, \cdot)$ does not necessarily fulfill the triangle inequality, and is also not symmetric. We therefore need the bounds of the following lemma.

► **Lemma 18.** *For any level $i \geq 1$ and vertices v and u that are cut by \mathcal{D} on level $i - 1$ we have $\text{dist}'(v, u) \leq \text{dist}(v, u) + 70 \cdot \rho 2^i$. Let $B \in \mathcal{B}_j$ be the part on some level $j \geq i$ with $v, u \in B$. For any $p \in I_B$ we have $\text{dist}'(v, u) \leq \text{dist}(v, u) + \langle \text{dist}(p, u) \rangle_j$.*

Proof. Let $B' \in \mathcal{B}_i$ be the part on level i containing both v and u . By Lemma 5 there is an interface point $p' \in I_{B'}$ such that $\text{dist}(v, p') + \text{dist}(p', u) \leq \text{dist}(v, u) + 34 \cdot \rho 2^i$. By definition of the rounding we also have $\langle \text{dist}(p', u) \rangle_i \leq \text{dist}(p', u) + 36 \cdot \rho 2^i$. Hence $\text{dist}'(v, u) \leq \text{dist}(v, p') + \langle \text{dist}(p', u) \rangle_i \leq \text{dist}(v, p') + \text{dist}(p', u) + 36 \cdot \rho 2^i \leq \text{dist}(v, u) + 70 \cdot \rho 2^i$.

The second part is obvious if $j = i$ from the definition of $\text{dist}'(v, u)$. If $j \geq i + 1$, we use the above bound on $\text{dist}'(v, u)$ together with the additive shift of the rounding and the triangle inequality of $\text{dist}(\cdot, \cdot)$ to obtain

$$\begin{aligned} \text{dist}'(v, u) &\leq \text{dist}(v, u) + 70 \cdot \rho 2^i \leq \text{dist}(v, p) + \text{dist}(p, u) + 70 \cdot \rho 2^{j-1} \\ &\leq \text{dist}(v, p) + \langle \text{dist}(p, u) \rangle_j - 35 \cdot \rho 2^j + 70 \cdot \rho 2^{j-1} = \text{dist}(v, p) + \langle \text{dist}(p, u) \rangle_j. \blacktriangleleft \end{aligned}$$

For any non-empty set S of facilities, we define $\text{dist}'(v, S) = \min_{f \in S} \{ \text{dist}'(v, f) \}$, and for empty sets we let $\text{dist}'(v, \emptyset) = \infty$. Analogous to $\text{cost}_{\mathcal{I}_0}(S)$, for a solution S to some instance \mathcal{I}_0 we also define $\text{cost}'_{\mathcal{I}_0}(S)$ using $\text{dist}'(\cdot, \cdot)$ as

$$\text{cost}'_{\mathcal{I}_0}(S) = \sum_{v \in V} \chi_{\mathcal{I}_0}(v) \cdot \text{dist}'(v, S)^q.$$

We show the following lemma, which translates between $\text{cost}'_{\mathcal{I}_D}$ and $\text{cost}_{\mathcal{I}}$, and is implied by the preciseness of the interface sets and the fact that \mathcal{I}_D has small distortion. Recall that the set of facilities is the same in \mathcal{I} and \mathcal{I}_D , i.e., a solution to one of these instances is also a solution to the other.

► **Lemma 19** (\star). *Let \mathcal{I} be an instance of k -CLUSTERING^q or FACILITY LOCATION^q with optimum solution OPT and approximate solution L . Let \mathcal{I}_D be an instance of small distortion for some $0 < \varepsilon < 1/2$, computed from L and a hierarchical decomposition \mathcal{D} with precise interface sets for $\rho \leq \frac{\varepsilon^{q+4+1/q}}{280\sigma(q+1)^q}$ according to Lemma 5. For the witness solution \hat{S} of \mathcal{I}_D we have $\text{cost}'_{\mathcal{I}_D}(\hat{S}) \leq (1 + O(\varepsilon)) \text{cost}_{\mathcal{I}}(OPT) + O(\varepsilon) \text{cost}_{\mathcal{I}}(L)$. Moreover, for any solution S we have $\text{cost}_{\mathcal{I}}(S) \leq (1 + O(\varepsilon)) \text{cost}'_{\mathcal{I}_D}(S) + O(\varepsilon) \text{cost}_{\mathcal{I}}(L)$.*

The next lemma states the properties of the dynamic program that for any coarse instance \mathcal{I}_0 computes an *optimal rounded interface-respecting solution*, which formally is a subset OPT' of facilities that minimizes $\text{cost}'_{\mathcal{I}_0}(OPT')$ with $|OPT'| \leq k$ for k -CLUSTERING^q, while for FACILITY LOCATION^q it minimizes $\text{cost}'_{\mathcal{I}_0}(OPT') + \sum_{f \in OPT'} w_f$. This step of the algorithm exploits the conciseness of the interface sets and the coarseness of the instance to bound the runtime. We prove the following lemma in Section 4.2.

► **Lemma 20.** *Let \mathcal{I}_0 be an instance of k -CLUSTERING^q or FACILITY LOCATION^q that for some $\varepsilon > 0$ is coarse w.r.t. a hierarchical decomposition \mathcal{D} with concise interface sets for some $1/2 \geq \rho > 0$ according to Lemma 5. An optimum rounded interface-respecting solution for \mathcal{I}_0 can be computed in $(nX/\varepsilon)^{(h/\rho)^{O(1)}}$ time.*

We are now ready to put together the above lemmas to prove Theorem 2. Due to space constraints however, the formal proof is deferred to the full version of the paper.

4.2 The dynamic program (proof of Lemma 20)

We describe the algorithm for k -CLUSTERING^q, and only mention in the end how to modify the algorithm to compute a solution for FACILITY LOCATION^q.

The solution is computed by a dynamic program recursing on the decomposition \mathcal{D} . Let W be the vertex set that \mathcal{D} decomposes, and which contains all clients and facilities of the coarse instance \mathcal{I} . Roughly speaking, the table of the dynamic program will have an entry for every part $B \in \mathcal{B}_i$ of \mathcal{D} on all levels $i \geq \xi(W)$, for which it will estimate the distance from each interface point on all higher levels $j \geq i + 1$ to the closest facility of the optimum solution. That is, if $\tilde{B} \in \mathcal{B}_j$ is a higher-level part for which $B \subseteq \tilde{B}$, then the distances from all interface points $I_{\tilde{B}}$ to facilities of the solution in \tilde{B} will be estimated.

Here the estimation happens in two ways. First off, the distances to facilities outside of B have to be guessed. That is, there is an *external* distance function d_j^+ that assigns a distance to each interface point of $I_{\tilde{B}}$, anticipating the distance from such a point to the closest facility of \tilde{B} , if this facility lies outside of B . In order to verify whether the guess was correct, each entry for a part B on level i also provides an *internal* distance function d_j^- , which stores the distance from each interface point of $I_{\tilde{B}}$ on level $j \geq i + 1$ to the closest facility, if the facility is guessed to lie inside of B .

The other way in which distances are estimated concerns the preciseness with which they are stored. The distance functions d_j^+ and d_j^- will only take rounded values $\langle x \rangle_j$ where $0 < x \leq 2^{j+5}$, or ∞ if no facility at the appropriate distance exists. In particular, if the facility of the solution in \tilde{B} that is closest to $p \in I_{\tilde{B}}$ lies outside of B then $d_j^-(p) = \infty$, and if it lies inside of B then $d_j^+(p) = \infty$. If there is no facility of the solution in \tilde{B} then both distance functions d_j^+ and d_j^- are set to ∞ for all $p \in I_{\tilde{B}}$. Note that this means that at least one of $d_j^+(p)$ and $d_j^-(p)$ is always set to ∞ . Note also that the finite values in the domains of the distance functions admit to store the rounded distance to any facility in \tilde{B} on level j , since the diameter of \tilde{B} is at most 2^{j+4} by Lemma 13, and the distance from any $p \in I_{\tilde{B}}$ to \tilde{B} is at most $(1 + 2\rho) \text{diam}(\tilde{B})$ by Lemma 16, i.e., for any $f \in \tilde{B} \cap F$ we have $\text{dist}(p, f) \leq (1 + 2\rho)2^{j+4} \leq 2^{j+5}$ using $\rho \leq 1/2$.

Formal definition of the table. Let us denote by $I_{\tilde{B}}^j$ the interface set of the part $\tilde{B} \in \mathcal{B}_j$ on level $j \geq i + 1$ containing $B \in \mathcal{B}_i$, i.e., $I_{\tilde{B}}^j = I_{\tilde{B}}$. Every entry of the dynamic programming table T is defined by a part $B \in \mathcal{B}_i$ of \mathcal{D} on a level $i \in \{\xi(W), \dots, \lambda(W)\}$, and two distance functions $d_j^+, d_j^- : I_{\tilde{B}}^j \rightarrow \{\langle x \rangle_j \mid 0 < x \leq 2^{j+5}\} \cup \{\infty\}$ for each $j \in \{i + 1, \dots, \lambda(W)\}$, such that $\max\{d_j^+(p), d_j^-(p)\} = \infty$ for all $p \in I_{\tilde{B}}^j$. Additionally, each entry comes with an integer $k' \in \{0, \dots, k\}$, which is a guess on the number of facilities that the optimum solution contains in B .

In an entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ we store the rounded interface-respecting cost of connecting the clients of B to facilities that adhere to the distance functions. More concretely, let $S \subseteq F \cap B$ be any subset of facilities in B . We say that S is *compatible* with an entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ if $|S| = k'$, and for any $j \geq i + 1$ the values of the distance functions for every interface point $p \in I_{\tilde{B}}^j$ are set to either

- $d_j^-(p) = \langle \text{dist}(p, S) \rangle_j$ and $d_j^+(p) = \infty$, or
- $d_j^+(p) \leq \langle \text{dist}(p, S) \rangle_j$ and $d_j^-(p) = \infty$.

Recall that $\text{dist}(v, \emptyset) = \infty$, and so the empty set $S = \emptyset$ is compatible with an entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ if $k' = 0$, and the values of all internal distance functions are set to ∞ . Over all sets $S \subseteq F \cap B$ compatible with the entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ for $B \in \mathcal{B}_i$, the entry should store the minimum value of $C_B(S)$, which is defined as

$$C_B(S) = \sum_{v \in B} \chi_{\mathcal{I}_0}(v) \cdot \min \left\{ \text{dist}'(v, S), \min_{\substack{j \geq i+1 \\ p \in I_B^j}} \{ \text{dist}(v, p) + d_j^+(p) \} \right\}.$$

If there is no compatible set $S \subseteq F \cap B$ for the entry, then $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}] = \infty$.

On the highest level $i = \lambda(W)$, there are no distance functions to adhere to on levels $j \geq i + 1$, and thus any set $S \subseteq W$ of facilities is compatible with the entry for $B = W$ and $k' = |S|$. Furthermore, $\text{cost}'_{\mathcal{I}_0}(S)$ is equal to $C_W(S)$, since W contains all clients and facilities of the coarse instance \mathcal{I}_0 . In particular, the entry of T for which $k' = k$ and $B = W$, will contain the objective function value of the optimum rounded interface-respecting solution to \mathcal{I}_0 . Hence if we can compute the table T we can also output the optimum rounded interface-respecting solution via this entry.

Computing the table. We begin with a part $B \in \mathcal{B}_{\xi(W)}$ on the lowest considered level $\xi(W)$, for which we know that B contains at most one facility, as \mathcal{I}_0 is coarse. If B contains no facility, then only $S = \emptyset$ can be compatible with the entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ and computing the value of the entry is straightforward given the definition of $C_B(S)$, where all incompatible entries are set to ∞ . If B contains one facility f , then any compatible set S is either empty or only contains f . We can thus check whether either of the two options is compatible with the entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ by checking if k' is set to 0 or 1, respectively, and checking that all values of the internal distance function are set correctly. Thereafter we can again use the definition of $C_B(S)$ to compute the values for both possible sets S and store them in the respective compatible entries. All incompatible entries are set to ∞ .

Now fix a part $B \in \mathcal{B}_i$ that lies on a level $i > \xi(W)$. We show how to compute all entries $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ for all values k' and distance functions. By induction we have already computed the correct values of all entries of T for parts $B' \in \mathcal{B}_{i-1}$ where $B' \subseteq B$. We order these parts arbitrarily, so that B'_1, \dots, B'_ℓ are the parts of \mathcal{B}_{i-1} contained in B . We then define an auxiliary table \hat{T} that is similar to the table T , but should compute the best compatible facility set in the union $B'_{\leq \ell} = \bigcup_{h=1}^{\ell} B'_h$ of the first ℓ subparts of B . Accordingly, \hat{T} has an entry for each union of parts $B'_{\leq \ell}$, each $k' \in \{0, \dots, k\}$, and distance functions $d_j^+, d_j^- : I_B^j \rightarrow \{\langle x \rangle_j \mid 0 < x \leq 2^{j+5}\} \cup \{\infty\}$ for each $j \in \{i, \dots, \lambda(W)\}$, such that $\max\{d_j^+(p), d_j^-(p)\} = \infty$ for all $p \in I_B^j$. Here, naturally, $I_B^i = I_B$, i.e., the entry also takes the interface set of B into account.

Analogous to before, a set $S \subseteq F \cap B'_{\leq \ell}$ of facilities in the union is *compatible* with an entry $\hat{T}[B'_{\leq \ell}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$ if $|S| = k'$, and for any $j \geq i$ the values of the distance functions for every interface point $p \in I_B^j$ are set to either

- $d_j^-(p) = \langle \text{dist}(p, S) \rangle_j$ and $d_j^+(p) = \infty$, or
- $d_j^+(p) \leq \langle \text{dist}(p, S) \rangle_j$ and $d_j^-(p) = \infty$.

Over all sets $S \subseteq F \cap B'_{\leq \ell}$ compatible with $\hat{T}[B'_{\leq \ell}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$, the entry should store the minimum value of $\hat{C}_{\leq \ell}(S)$, which is defined as

$$\hat{C}_{\leq \ell}(S) = \sum_{v \in B'_{\leq \ell}} \chi_{\mathcal{I}_0}(v) \cdot \min \left\{ \text{dist}'(v, S), \min_{\substack{j \geq i \\ p \in I_B^j}} \{ \text{dist}(v, p) + d_j^+(p) \} \right\}.$$

If there is no compatible set $S \subseteq F \cap B'_{\leq \ell}$ for the entry, then $\hat{T}[B'_{\leq \ell}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}] = \infty$.

To compute T using the auxiliary table \hat{T} , note that since $B = B'_{\leq b}$, any set $S \subseteq F \cap B$ is compatible with the entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ if and only if it is compatible with a corresponding entry $\hat{T}[B'_{\leq b}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$ for some internal distance function d_i^- on level i . Furthermore, if $d_i^+(p) = \infty$ for all $p \in I^i$, then $C_B(S) = \hat{C}_{\leq b}(S)$ for such a set S . Therefore we can easily compute the entry $T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ from \hat{T} by setting

$$T[B, k', (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}] = \min_{d_i^-} \left\{ \hat{T}[B'_{\leq b}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}] \mid \forall p \in I_B^i : d_i^+(p) = \infty \right\}.$$

Computing the auxiliary table. Also computing an entry of \hat{T} for $B'_{\leq 1}$ is easy using the entries of T for B'_1 , since $B'_1 = B'_{\leq 1}$ and so (taking the index shift of i into account) we have

$$\hat{T}[B'_{\leq 1}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}] = T[B'_1, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}].$$

To compute entries of \hat{T} for some $B'_{\leq \ell}$ where $\ell \geq 2$, we combine entries of table T for B'_ℓ with entries of table \hat{T} for $B'_{\leq \ell-1}$. However we will only combine entries with distance functions that imply compatible solutions. More concretely, we say that distance functions $(d_j^+, d_j^-)_{j=i}^{\lambda(W)}$ for $B'_{\leq \ell}$, $(\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}$ for B'_ℓ , and $(\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}$ for $B'_{\leq \ell-1}$ are *consistent* if for every level $j \geq i$ and $p \in I_B^j$ we have one of

1. $d_j^+(p) = \delta_j^+(p) = \beta_j^+(p)$ and $d_j^-(p) = \delta_j^-(p) = \beta_j^-(p) = \infty$, or
2. $d_j^-(p) = \delta_j^-(p) = \beta_j^-(p)$ and $d_j^+(p) = \delta_j^+(p) = \beta_j^+(p) = \infty$, or
3. $d_j^+(p) = \delta_j^+(p) = \beta_j^+(p)$ and $d_j^-(p) = \delta_j^-(p) = \beta_j^-(p) = \infty$.

The algorithm now considers all sets of consistent distance functions to compute an entry $\hat{T}[B'_{\leq \ell}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$ for $\ell \geq 2$ by setting it to

$$\min \left\{ T[B'_\ell, k'', (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}] + \hat{T}[B'_{\leq \ell-1}, k' - k'', (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}] \mid k'' \in \{0, \dots, k'\} \text{ and } (d_j^+, d_j^-)_{j=i}^{\lambda(W)}, (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)} \text{ are consistent} \right\} \quad (1)$$

We now prove the correctness using two lemmas. The following lemma implies that if we only consider consistent distance functions to compute entries recursively, then the entries will store values for compatible solutions.

► **Lemma 21.** *Let $(d_j^+, d_j^-)_{j=i}^{\lambda(W)}$ for $B'_{\leq \ell}$, $(\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}$ for B'_ℓ , and $(\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}$ for $B'_{\leq \ell-1}$ be consistent distance functions, and let $S_1 = B'_\ell \cap F$ and $S_2 = B'_{\leq \ell-1} \cap F$ be facility sets. If S_1 is compatible with entry $T[B'_\ell, |S_1|, (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}]$ and S_2 is compatible with entry $\hat{T}[B'_{\leq \ell-1}, |S_2|, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}]$, then the union $S = S_1 \cup S_2$ is compatible with entry $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$. Moreover, $\hat{C}_{\leq \ell}(S) = C_{B'_\ell}(S_1) + \hat{C}_{\leq \ell-1}(S_2)$.*

Proof. To prove compatibility of S with the entry $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$, it suffices to show that the distance functions are set correctly. Fix a level $j \geq i$ and an interface point $p \in I_B^j$. There are three cases to consider, according to the definition of consistency of the distance functions. In the first case, all three internal distance functions are set to ∞ , and all external distance functions are set to the same value. In particular, since S_1 and S_2 are compatible with their respective entries, we have $d_j^+(p) = \delta_j^+(p) = \beta_j^+(p) \leq \min\{\langle \text{dist}(p, S_1) \rangle_j, \langle \text{dist}(p, S_2) \rangle_j\} = \langle \text{dist}(p, S) \rangle_j$, as $S = S_1 \cup S_2$. In the second case, $\beta_j^-(p) = \delta_j^-(p) = \infty$ and so $\beta_j^+(p) \leq \langle \text{dist}(p, S_2) \rangle_j$, since S_2 is compatible with its entry, and

$\delta_j^-(p) = \langle \text{dist}(p, S_1) \rangle_j$, since S_1 is compatible with its entry. Since we also have $\beta_j^+(p) = \delta_j^-(p)$ we get $\langle \text{dist}(p, S_1) \rangle_j \leq \langle \text{dist}(p, S_2) \rangle_j$, and hence $\langle \text{dist}(p, S) \rangle_j = \langle \text{dist}(p, S_1) \rangle_j$. Consistency furthermore implies $d_j^-(p) = \delta_j^-(p) = \langle \text{dist}(p, S) \rangle_j$ and $d_j^+(p) = \infty$. The third case is analogous to the second, and therefore S is compatible with its entry.

For the second part, we consider the contributions of vertices to the terms $\hat{C}_{\leq \ell}(S)$, $C_{B'_\ell}(S_1)$, and $\hat{C}_{\leq \ell-1}(S_2)$, and show that they are the same for $\hat{C}_{\leq \ell}(S)$ and for $C_{B'_\ell}(S_1) + \hat{C}_{\leq \ell-1}(S_2)$. For this we first fix a vertex $v \in B'_{\leq \ell-1}$, and in the following distinguish the cases where its contribution to $\hat{C}_{\leq \ell-1}(S_2)$ and $\hat{C}_{\leq \ell}(S)$ is due to a facility or an interface point.

The first case is that $\text{dist}'(v, S_2) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, i.e., the contribution of v to $\hat{C}_{\leq \ell-1}(S_2)$ is given by a facility of S_2 . Note that the consistency of the distance functions always implies that $\beta_j^+(p) = d_j^+(p)$ or $d_j^+(p) = \infty$ for any level $j \geq i$ and interface point $p \in I_B^j$, and so $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$. At the same time $\text{dist}'(v, S) \leq \text{dist}'(v, S_2)$ as $S_2 \subseteq S$. We hence get that $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$, i.e., the contribution of v to $\hat{C}_{\leq \ell}(S)$ is also given by a facility of S in this case. Thus to show that the contribution of v to $\hat{C}_{\leq \ell-1}(S_2)$ and $\hat{C}_{\leq \ell}(S)$ is the same, we need to show that $\text{dist}'(v, S) = \text{dist}'(v, S_2)$. Note that this is implied if $\text{dist}'(v, S) \geq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, since we have $\text{dist}'(v, S) \leq \text{dist}'(v, S_2) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$. Thus the following proves the claim, using that the contribution of v to $\hat{C}_{\leq \ell}(S)$ is given by a facility of S .

▷ **Claim 22.** For $v \in B'_{\leq \ell-1}$, if $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$ then we have $\text{dist}'(v, S) = \text{dist}'(v, S_2)$ or $\text{dist}'(v, S) \geq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$.

Proof. Given $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$, assume to the contrary that we have $\text{dist}'(v, S) < \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$ and $\text{dist}'(v, S) \neq \text{dist}'(v, S_2)$, which, as $S = S_1 \cup S_2$, means $\text{dist}'(v, S) < \text{dist}'(v, S_2)$. The latter inequality implies that the value of $\text{dist}'(v, S)$ is obtained for some facility $f \in S_1 \subseteq B'_\ell$. In particular, $v \in B'_{\leq \ell-1}$ and $f \in B'_\ell$ are cut at level $i - 1$, and so there is an interface point $p \in I_B^i$ such that $\text{dist}'(v, S) = \text{dist}(v, p) + \langle \text{dist}(p, f) \rangle_i$, and f is the closest facility to p in S , i.e. $\langle \text{dist}(p, S) \rangle_i = \langle \text{dist}(p, f) \rangle_i$. Using the former of the assumed inequalities we get $\text{dist}(v, p) + \langle \text{dist}(p, f) \rangle_i = \text{dist}'(v, S) < \text{dist}(v, p) + \beta_i^+(p)$, and so we can conclude that $\langle \text{dist}(p, f) \rangle_i < \beta_i^+(p)$.

Using the inequality of the premise of the claim, we also get $\text{dist}(v, p) + \langle \text{dist}(p, f) \rangle_i = \text{dist}'(v, S) \leq \text{dist}(v, p) + d_i^+(p)$, i.e. $\langle \text{dist}(p, f) \rangle_i \leq d_i^+(p)$. Since S is compatible with entry $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$, we have $d_i^+(p) = \infty$ or $d_i^+(p) \leq \langle \text{dist}(p, S) \rangle_i$. In the latter case we would have $d_i^+(p) \leq \langle \text{dist}(p, S) \rangle_i = \langle \text{dist}(p, f) \rangle_i < \beta_i^+(p)$, which however cannot happen if the distance functions are consistent. Thus compatibility of S implies $d_i^+(p) = \infty$ and $d_i^-(p) = \langle \text{dist}(p, f) \rangle_i$. In particular, we can conclude that $d_i^-(p)$ has a finite value (as f exists) and $\beta_i^+(p)$ differs from $d_i^-(p)$. This can only mean that the third of the consistency properties applies to p at level i , and so $\beta_i^-(p) = d_i^-(p) = \langle \text{dist}(p, f) \rangle_i$.

In particular, also $\beta_i^-(p)$ has a finite value, and using the compatibility of S_2 with entry $\hat{T}[B'_{\leq \ell-1}, |S_2|, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}]$, we can conclude that there exists a facility $f' \in S_2 \subseteq B'_{\leq \ell-1}$ with $\langle \text{dist}(p, f') \rangle_i = \beta_i^-(p) = \langle \text{dist}(p, f) \rangle_i$. Now let $j \leq i$ be the level for which $v \in B'_{\leq \ell-1}$ and $f' \in B'_{\leq \ell-1}$ are cut at level $j - 1$ by \mathcal{D} . Lemma 18 implies $\text{dist}'(v, f') \leq \text{dist}(v, p) + \langle \text{dist}(p, f') \rangle_i$, but then we have

$$\text{dist}'(v, S_2) \leq \text{dist}'(v, f') \leq \text{dist}(v, p) + \langle \text{dist}(p, f') \rangle_i = \text{dist}(v, p) + \langle \text{dist}(p, f) \rangle_i = \text{dist}'(v, S),$$

which is a contradiction to $\text{dist}'(v, S) < \text{dist}'(v, S_2)$. \triangleleft

The next case we consider is that $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\} < \text{dist}'(v, S)$, i.e., the contribution of v to $\hat{C}_\ell(S)$ is given by an interface point. As observed before, we have $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$ and $\text{dist}'(v, S) \leq \text{dist}'(v, S_2)$, which implies $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} < \text{dist}'(v, S_2)$, i.e. in this case the contribution of v to $\hat{C}_{\ell-1}(S_2)$ is also given by an interface point. Note that it also implies $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, and thus the following claim shows that the contribution of v to $\hat{C}_\ell(S)$ and $\hat{C}_{\ell-1}(S_2)$ is the same.

▷ **Claim 23.** For $v \in B'_{\leq \ell-1}$, if $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$ then we have $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} = \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$.

Proof. Given $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, assume to the contrary that $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} \neq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$. As observed before, the consistency of the distance functions always implies $\beta_j^+(p) = d_j^+(p)$ or $d_j^+(p) = \infty$, and thus we must have $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} < \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$. Let $j \geq i$ and $p \in I_B^j$ be the level and interface point for which the minimum of the former term of this inequality is obtained. The inequality then implies $\beta_j^+(p) < d_j^+(p)$ for this particular point p and level j , which can only be the case if $\beta_j^+(p) < \infty$ and $d_j^+(p) = \infty$. The values of $\beta_j^+(p)$ and $d_j^+(p)$ can only differ if the second of the consistency properties applies to p at level j , and so $\beta_j^+(p) = d_j^-(p)$. Since $\beta_j^+(p) < \infty$, the compatibility of S with entry $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$, implies $\beta_j^+(p) = d_j^-(p) = \langle \text{dist}(p, S) \rangle_j$.

Now let $f \in S \subseteq B'_{\leq \ell}$ be the facility for which $\text{dist}'(v, S) = \text{dist}'(v, f)$ (which exists as $d_j^-(p) < \infty$). Let $j' \leq i$ be the level for which $v \in B'_{\leq \ell-1}$ and $f \in B'_{\leq \ell}$ are cut at level $j' - 1$ by \mathcal{D} . By Lemma 18 we have $\text{dist}'(v, f) \leq \text{dist}(v, p) + \langle \text{dist}(p, f) \rangle_j$, since $j' \leq j$ and the part $B \in \mathcal{B}_i$ containing v and f is itself contained in some part $\tilde{B} \in \mathcal{B}_j$ with $v, f \in \tilde{B}$ and $p \in I_{\tilde{B}}$. But then,

$$\text{dist}'(v, S) \leq \text{dist}(v, p) + \langle \text{dist}(p, f) \rangle_j = \text{dist}(v, p) + \langle \text{dist}(p, S) \rangle_j = \text{dist}(v, p) + \beta_j^+(p).$$

However the last term is equal to $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, which gives a contradiction to our premise $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$. ◁

So far we considered the case when the contribution of v to $\hat{C}_{\ell-1}(S_2)$ is given by a facility, or when the contribution of v to $\hat{C}_\ell(S)$ is given by an interface point. Thus the last case we consider is when the contribution of v to $\hat{C}_{\ell-1}(S_2)$ is given by an interface point, and the contribution of v to $\hat{C}_\ell(S)$ is given by a facility, i.e., $\min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\} < \text{dist}'(v, S_2)$ and $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$. We need to show that $\text{dist}'(v, S) = \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$. First assume $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$. Due to Claim 23 this would imply $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$, which however contradicts our assumption to the contrary, i.e., that the contribution of v to $\hat{C}_\ell(S)$ is given by a facility. Hence we must instead have $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$.

According to Claim 22, our assumption that $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + d_j^+(p)\}$ implies $\text{dist}'(v, S) = \text{dist}'(v, S_2)$ or $\text{dist}'(v, S) \geq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$. In the former case, together with our assumption that the contribution of v to $\hat{C}_{\ell-1}(S_2)$ is given by an interface point, we would get $\text{dist}'(v, S) > \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, for which we saw above that this leads to a contradiction via Claim 23. Hence we are left with the other

implication of Claim 22, i.e., $\text{dist}'(v, S) \geq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$. This together with our conclusion from above, i.e., $\text{dist}'(v, S) \leq \min_{j \geq i, p \in I_B^j} \{\text{dist}(v, p) + \beta_j^+(p)\}$, means that the contribution of v to $\hat{C}_\ell(S)$ and $\hat{C}_{\ell-1}(S_2)$ is the same.

By analogous arguments, the contribution of any $v \in B'_\ell$ to $C_{B'_\ell}(S_1)$ is the same as its contribution to $\hat{C}_{\leq \ell}(S)$. Since B'_ℓ and $B'_{\leq \ell-1}$ partition the set $B'_{\leq \ell}$, this means that $\hat{C}_{\leq \ell}(S) = C_{B'_\ell}(S_1) + \hat{C}_{\leq \ell-1}(S_2)$, as required. \blacktriangleleft

The next lemma implies that the compatible facility set minimizing $\hat{C}_{\leq \ell}(S)$ is considered as a solution when recursing over consistent distance functions.

► **Lemma 24** (\star). *Let $S = B'_{\leq \ell} \cap F$ be a facility set of $B'_{\leq \ell}$ that is compatible with entry $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$, and let $S_1 = S \cap B'_\ell$ and $S_2 = S \cap B'_{\leq \ell-1}$. Then there exist distance functions $(\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}$ for B'_ℓ , and $(\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}$ for $B'_{\leq \ell-1}$ such that*

- $(d_j^+, d_j^-)_{j=i}^{\lambda(W)}$, $(\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}$, and $(\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}$ are consistent, and
- the set S_1 is compatible with entry $T[B'_\ell, |S_1|, (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}]$ and S_2 is compatible with entry $\hat{T}[B'_{\leq \ell-1}, |S_2|, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}]$.

To argue that the algorithm sets the value of $\hat{T}[B'_{\leq \ell}, k', (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$ correctly via (1), consider a set $S \subseteq B'_{\leq \ell}$ that is compatible with this entry and minimizes $\hat{C}_{\leq \ell}(S)$. By induction, Lemma 24 implies $T[B'_\ell, |S_1|, (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}] \leq C_{B'_\ell}(S_1)$ and $\hat{T}[B'_{\leq \ell-1}, |S_2|, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}] \leq \hat{C}_{\leq \ell-1}(S_2)$, where $S_1 = S \cap B'_\ell$ and $S_2 = S \cap B'_{\leq \ell-1}$. From (1) we therefore obtain $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}] \leq C_{B'_\ell}(S_1) + \hat{C}_{\leq \ell-1}(S_2)$. By Lemma 21 only compatible sets are stored in an entry by induction, and so the definition of S implies $\hat{T}[B'_{\leq \ell}, |S|, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}] = \hat{C}_{\leq \ell}(S)$, as required.

Bounding the runtime. To bound the size of the tables T and \hat{T} , note that since there are $\lambda(W) - \xi(W) + 1 \leq 2 \log_2(nX/\varepsilon) + 2$ considered levels i , and each level \mathcal{B}_i of \mathcal{D} is a partition of W where $|W| \leq n$, there are at most $O(n \log(nX/\varepsilon))$ parts B considered by T in total. The other table \hat{T} considers the same number of parts, since a set $B'_{\leq \ell}$ can be uniquely mapped to the part B'_ℓ . The number of possible values for k' is $k + 1 = O(n)$. The domain $\{\langle x \rangle_j \mid 0 < x \leq 2^{j+5}\} \cup \{\infty\}$ of a distance function for level j has at most $\lceil 2^{j+5}/(\rho 2^j) \rceil + 1 = O(1/\rho)$ values, since $\langle x \rangle_j$ rounds a value to a multiple of $\rho 2^j$. The conciseness of the interface sets means that $|I_B^j| \leq (h/\rho)^{O(1)}$ according to Lemma 5. Hence there are at most $O(1/\rho)^{(h/\rho)^{O(1)}} = 2^{(h/\rho)^{O(1)}}$ possible distance functions. Since each entry of the table stores two distance functions for each of at most $2 \log_2(nX/\varepsilon) + 2$ levels, the total number of entries of T and \hat{T} is at most

$$O(n \log(nX/\varepsilon)) \cdot n \cdot (2^{(h/\rho)^{O(1)}})^{O(\log(nX/\varepsilon))} = (nX/\varepsilon)^{(h/\rho)^{O(1)}}.$$

Computing an entry of a table is dominated by (1). Going through all values $k' \leq n$ and all possible consistent distance functions to compute (1), takes $n \cdot 2^{(h/\rho)^{O(1)}}$ time, as there are $2^{(h/\rho)^{O(1)}}$ possible distance functions. Hence the total runtime is $(nX/\varepsilon)^{(h/\rho)^{O(1)}}$, proving Lemma 20.

The Facility Location^q problem. To compute an optimum rounded interface-respecting solution to FACILITY LOCATION^q, the tables T and \hat{T} can ignore the number of open facilities k' , i.e., they have respective entries $T[B, (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ and $\hat{T}[B'_{\leq \ell}, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$.

Accordingly, compatibility of facility sets with entries is defined as before, but ignoring the sizes of the sets. The value stored in each entry now also takes the opening costs of facilities into account. That is, for any set of facilities $S \subseteq F \cap B$ in a part B we define

$$C_B(S) = \sum_{v \in B} \chi_{\mathcal{I}_0}(v) \cdot \min \left\{ \text{dist}'(v, S), \min_{\substack{j \geq i+1 \\ p \in I_B^j}} \left\{ \text{dist}(v, p) + d_j^+(p) \right\} \right\} + \sum_{f \in S} w_f,$$

and an entry $T[B, (d_j^+, d_j^-)_{j=i+1}^{\lambda(W)}]$ stores the minimum value of $C_B(S)$ over all sets S compatible with the entry, or ∞ if no such set exists. For $S \subseteq F \cap B'_{\leq \ell}$ in a union of subparts $B'_{\leq \ell}$ we define

$$\hat{C}_{\leq \ell}(S) = \sum_{v \in B'_{\leq \ell}} \chi_{\mathcal{I}_0}(v) \cdot \min \left\{ \text{dist}'(v, S), \min_{\substack{j \geq i \\ p \in I_B^j}} \left\{ \text{dist}(v, p) + d_j^+(p) \right\} \right\} + \sum_{f \in S} w_f,$$

and an entry $\hat{T}[B'_{\leq \ell}, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}]$ stores the minimum value of $\hat{C}_{\leq \ell}(S)$ over all sets S compatible with the entry, or ∞ if no such set exists.

The entries of the tables can be computed in the same manner as before, but ingoring the sets sizes. In particular, the most involved recursion becomes

$$\begin{aligned} \hat{T}[B'_{\leq \ell}, (d_j^+, d_j^-)_{j=i}^{\lambda(W)}] = \min \{ & T[B'_\ell, (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}] + \hat{T}[B'_{\leq \ell-1}, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)}] \mid \\ & (d_j^+, d_j^-)_{j=i}^{\lambda(W)}, (\delta_j^+, \delta_j^-)_{j=i}^{\lambda(W)}, (\beta_j^+, \beta_j^-)_{j=i}^{\lambda(W)} \text{ are consistent} \}. \end{aligned}$$

Note that if $S_1 = B'_\ell \cap F$ and $S_2 = B'_{\leq \ell-1} \cap F$ then these two sets are disjoint, and so $\sum_{f \in S} w_f = \sum_{f \in S_1} w_f + \sum_{f \in S_2} w_f$ for the union $S = S_1 \cup S_2$. Hence when proving $\hat{C}_{\leq \ell}(S) = C_{B'_\ell}(S_1) + \hat{C}_{\leq \ell-1}(S_2)$ for Lemma 21, we can ignore the facility opening costs, and the proof remains the same as before. All other arguments carry over, and thus an optimum rounded interface-respecting solution for an instance of FACILITY LOCATION^q can also be computed in $(nX/\varepsilon)^{(h/\rho)^{O(1)}}$ time.

5 Hardness for graphs of highway dimension 1

For both k -CLUSTERING^q and FACILITY LOCATION^q we present the same reduction from the NP-hard satisfiability problem (SAT), in which a boolean formula φ in conjunctive normal form is given, and a satisfying assignment of its variables needs to be found.

For a given SAT formula φ with k variables and ℓ clauses we construct a graph G_φ as follows. For each variable x we introduce a path $P_x = (t_x, u_x, f_x)$ with two edges of length 1 each. The two endpoints t_x and f_x are facilities of F and the additional vertex u_x is a client, i.e., $\chi(u_x) = 1$. For each clause C_i , where $i \in [\ell]$, we introduce a vertex v_i and add the edge $v_i t_x$ for each variable x such that C_i contains x as a positive literal, and we add the edge $v_i f_x$ for each x for which C_i contains x as a negative literal. Every edge incident to v_i has length $(11c)^i$ for the constant $c > 4$ due to Definition 1, and v_i is also a client, i.e., $\chi(v_i) = 1$. In case of FACILITY LOCATION^q, every facility $f \in F$ has cost $w_f = 1$, i.e., we construct an instance of the uniform version of the problem.

► **Lemma 25.** *The constructed graph G_φ has highway dimension 1.*

Proof. Fix a scale $r > 0$ and let $i = \lfloor \log_{11c}(r/5) + 1 \rfloor$. Note that $\beta_w(cr)$ cannot contain any edge incident to a vertex v_j for $j \geq i + 1$, since the length of every such edge is $(11c)^j \geq 11cr/5 > 2cr$ and the diameter of $\beta_w(cr)$ is at most $2cr$. Thus if $\beta_w(cr)$ contains a vertex v_j for $j \geq i + 1$, then $\beta_w(cr)$ contains only v_j , and there is nothing to prove. Note

also that any path in $\beta_w(cr)$ that does not use v_i has length at most $2 + \sum_{j=1}^{i-1} (2(11c)^j + 2)$, since any such path can contain at most two edges incident to a vertex v_j and the paths P_x of length 2 are connected only through edges incident to vertices v_j . The length of such a path is thus strictly shorter than

$$2 + 2 \left(\frac{(11c)^i}{11c - 1} - 1 \right) + 2i \leq 5(11c)^{i-1} \leq r,$$

where the first inequality holds since $i \geq 1$ and $c > 4$. Hence the only paths that need to be hit by hubs on scale r are those passing through v_i , which can clearly be done using only one hub, namely v_i . ◀

To finish the reduction for k -CLUSTERING^q, we claim that there is a satisfying assignment for φ if and only if there is a solution for G_φ with cost at most $k + \sum_{i=1}^{\ell} (11c)^{iq}$. If there is a satisfying assignment for φ we open each facility t_x for variables x that are set to true, and we open each facility f_x for variables x that are set to false. This opens exactly k facilities and the cost of the solution is $k + \sum_{i=1}^{\ell} (11c)^{iq}$, since each of the k vertices u_x is assigned to either t_x or f_x at distance 1, and vertex v_i is assigned to a vertex t_x or f_x at distance $(11c)^i$ that corresponds to a literal of C_i that is true.

Conversely, assume there is a solution to k -CLUSTERING^q of cost at most $k + \sum_{i=1}^{\ell} (11c)^{iq}$ in G_φ . Note that the minimum distance from any u_x to a facility is 1, while the minimum distance from any v_i to a facility is $(11c)^i$. Thus any solution must have cost at least $k + \sum_{i=1}^{\ell} (11c)^{iq}$, so that the assumed solution must open a facility at minimum distance for each client of G_φ . In particular, for each variable x , at least one of the facilities t_x and f_x is opened by the solution. Moreover, as only k facilities can be opened and there are k variables, exactly one of t_x and f_x is opened for each x . Thus the k -CLUSTERING^q solution in G_φ can be interpreted as an assignment for φ , where we set a variable x to true if t_x is opened, and we set it to false if f_x is opened. Since also for each v_i the solution opens a facility at minimum distance, there must be a variable in C_i that is set so that its literal in C_i is true, i.e., the assignment satisfies φ . Thus due to the above lemma bounding the highway dimension of G_φ , we obtain the Theorem 3 for k -CLUSTERING^q.

For FACILITY LOCATION^q we claim that there is a satisfying assignment for φ if and only if there is a solution for G_φ of cost at most $2k + \sum_{i=1}^{\ell} (11c)^{iq}$. In fact the arguments are exactly the same as for k -CLUSTERING^q above: if there is a satisfying assignment then a solution for FACILITY LOCATION^q of cost $2k + \sum_{i=1}^{\ell} (11c)^{iq}$ exists, by opening the k facilities corresponding to the assignment of cost 1 each. Conversely, any solution has cost at least $k + \sum_{i=1}^{\ell} (11c)^{iq}$ due to the edge lengths, and at least k facilities need to be opened, one for each variable gadget. This gives a minimum cost of $2k + \sum_{i=1}^{\ell} (11c)^{iq}$, and any such solution corresponds to a satisfying assignment of φ . This proves Theorem 3 for uniform FACILITY LOCATION^q.

References

- 1 I. Abraham, A. Fiat, A. V. Goldberg, and R. F. Werneck. Highway dimension, shortest paths, and provably efficient algorithms. In *SODA*, pages 782–793, 2010.
- 2 Sara Ahmadian, Ashkan Norouzi-Fard, Ola Svensson, and Justin Ward. Better guarantees for k-means and euclidean k-median by primal-dual algorithms. *SIAM Journal on Computing*, pages FOCS17–97, 2019.
- 3 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 '98, pages 106–113, New York, NY, USA, 1998. ACM. doi:10.1145/276698.276718.

- 4 Pranjali Awasthi, Moses Charikar, Ravishankar Krishnaswamy, and Ali Kemal Sinop. The hardness of approximation of euclidean k -means. In *31st International Symposium on Computational Geometry, SoCG 2015, June 22-25, 2015, Eindhoven, The Netherlands*, pages 754–767, 2015.
- 5 Yair Bartal and Lee-Ad Gottlieb. A linear time approximation scheme for euclidean TSP. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS*, 2013.
- 6 A. Becker, P. N. Klein, and D. Saulpic. Polynomial-time approximation schemes for k -center and bounded-capacity vehicle routing in metrics with bounded highway dimension. In *ESA*, pages 8:1–8:15, 2018.
- 7 J. Blum. Hierarchy of transportation network parameters and hardness results. In *IPEC*, pages 4:1–4:15, 2019.
- 8 Vladimir Braverman, Shaofeng H-C Jiang, Robert Krauthgamer, and Xuan Wu. Coresets for clustering in excluded-minor graphs and beyond. *arXiv preprint*, 2020. [arXiv:2004.07718](https://arxiv.org/abs/2004.07718).
- 9 Jaroslav Byrka, Thomas Penschel, Bartosz Rybicki, Srinivasan Aravind, and Khoa Trinh. An improved approximation for k -median, and positive correlation in budgeted optimization. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 737–756, 2015.
- 10 Vincent Cohen-Addad, Andreas Emil Feldmann, and David Saulpic. Near-linear time approximations schemes for clustering in doubling metrics. In *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pages 540–559, 2019. [doi:10.1109/FOCS.2019.00041](https://doi.org/10.1109/FOCS.2019.00041).
- 11 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. *SIAM Journal on Computing*, 48(2):644–667, 2019.
- 12 Y. Disser, A. E. Feldmann, M. Klimm, and J. Könemann. Travelling on graphs with small highway dimension. In *Graph-Theoretic Concepts in Computer Science - 45th International Workshop, WG*, volume 11789, pages 175–189. Springer, 2019.
- 13 A. E. Feldmann. Fixed-parameter approximations for k -center problems in low highway dimension graphs. *Algorithmica*, 81(3):1031–1052, 2019.
- 14 A. E. Feldmann, W. S. Fung, J. Könemann, and I. Post. A $(1 + \epsilon)$ -embedding of low highway dimension graphs into bounded treewidth graphs. *SIAM Journal on Computing*, 47(4):1275–1734, 2018.
- 15 Zachary Friggstad, Mohsen Rezapour, and Mohammad R. Salavatipour. Local search yields a PTAS for k -means in doubling metrics. *SIAM J. Comput.*, 48(2):452–480, 2019. [doi:10.1137/17M1127181](https://doi.org/10.1137/17M1127181).
- 16 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 '03*, 2003.
- 17 Anupam Gupta and Kanat Tangwongsan. Simpler analyses of local search algorithms for facility location. *CoRR*, abs/0809.2554, 2008. [arXiv:0809.2554](https://arxiv.org/abs/0809.2554).
- 18 Stavros G Kolliopoulos and Satish Rao. A nearly linear-time approximation scheme for the euclidean k -median problem. *SIAM Journal on Computing*, 37(3):757–782, 2007.
- 19 Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Inf. Comput.*, 222:45–58, 2013.
- 20 K. Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 281–290. ACM, 2004. [doi:10.1145/1007352.1007399](https://doi.org/10.1145/1007352.1007399).