



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/200960/>

Version: Published Version

Proceedings Paper:

Feldmann, A.E. and Marx, D. (2018) The parameterized hardness of the k-center problem in transportation networks. In: Eppstein, D., (ed.) Leibniz International Proceedings in Informatics, LIPIcs. 16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018), 18-20 Jun 2018, Malmö, Sweden. Schloss Dagstuhl--Leibniz-Zentrum für Informatik, 19:1-19:13. ISBN: 9783959770682. ISSN: 1868-8969.

<https://doi.org/10.4230/LIPIcs.SWAT.2018.19>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

The Parameterized Hardness of the k -Center Problem in Transportation Networks

Andreas Emil Feldmann¹

Department of Applied Mathematics, Charles University, Prague, Czechia
feldmann.a.e@gmail.com

Dániel Marx²

Institute for Computer Science and Control, Hungarian Academy of Sciences (MTA SZTAKI),
Budapest, Hungary
dmarx@cs.bme.hu

Abstract

In this paper we study the hardness of the k -CENTER problem on inputs that model transportation networks. For the problem, an edge-weighted graph $G = (V, E)$ and an integer k are given and a center set $C \subseteq V$ needs to be chosen such that $|C| \leq k$. The aim is to minimize the maximum distance of any vertex in the graph to the closest center. This problem arises in many applications of logistics, and thus it is natural to consider inputs that model transportation networks. Such inputs are often assumed to be planar graphs, low doubling metrics, or bounded highway dimension graphs. For each of these models, parameterized approximation algorithms have been shown to exist. We complement these results by proving that the k -CENTER problem is $W[1]$ -hard on planar graphs of constant doubling dimension, where the parameter is the combination of the number of centers k , the highway dimension h , and even the treewidth t . Moreover, under the Exponential Time Hypothesis there is no $f(k, t, h) \cdot n^{o(t + \sqrt{k+h})}$ time algorithm for any computable function f . Thus it is unlikely that the optimum solution to k -CENTER can be found efficiently, even when assuming that the input graph abides to all of the above models for transportation networks at once!

Additionally we give a simple parameterized $(1 + \varepsilon)$ -approximation algorithm for inputs of doubling dimension d with runtime $(k^k / \varepsilon^{O(kd)}) \cdot n^{O(1)}$. This generalizes a previous result, which considered inputs in D -dimensional L_q metrics.

2012 ACM Subject Classification Theory of computation \rightarrow Fixed parameter tractability, Theory of computation \rightarrow Facility location and clustering, Theory of computation \rightarrow Problems, reductions and completeness

Keywords and phrases k -center, parameterized complexity, planar graphs, doubling dimension, highway dimension, treewidth

Digital Object Identifier 10.4230/LIPIcs.SWAT.2018.19

Related Version <https://arxiv.org/abs/1802.08563>

1 Introduction

Given a graph $G = (V, E)$ with positive edge lengths $\ell : E \rightarrow \mathbb{Q}^+$, the k -CENTER problem asks to find k center vertices such that every vertex of the graph is as close as possible to one of centers. More precisely, if $\text{dist}(u, v)$ denotes the length of the shortest path between u

¹ Supported by project CE-ITI (GAČR no. P202/12/G061) of the Czech Science Foundation.

² Supported by ERC Consolidator Grant SYSTEMATICGRAPH (No. 725978)



© Andreas E. Feldmann and Daniel Marx;

licensed under Creative Commons License CC-BY

16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018).

Editor: David Eppstein; Article No. 19; pp. 19:1–19:13

Leibniz International Proceedings in Informatics



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

and v according to the edge lengths ℓ , let $B_v(\rho) = \{u \in V \mid \text{dist}(u, v) \leq \rho\}$ be the *ball of radius ρ around v* . A solution to k -CENTER is a set $C \subseteq V$ of *centers* such that $|C| \leq k$, and the objective is to minimize the *cost* of the solution, which is the smallest value ρ for which $\bigcup_{v \in C} B_v(\rho) = V$. This problem has numerous applications in logistics where easily accessible locations need to be chosen on a map under a budget constraint. For instance, a budget may be available to build k hospitals, shopping malls, or warehouses. These should be placed so that the distance from each point on the map to the closest facility is minimized.

The k -CENTER problem is NP-hard [28], and so *approximation algorithms* [28, 29] as well as *parameterized algorithms* [8, 11] have been developed for this problem. The former are algorithms that use polynomial time to compute an α -*approximation*, i.e., a solution that is at most α times worse than the optimum. For the latter, a *parameter p* is given as part of the input, and an optimum solution is computed in $f(p) \cdot n^{O(1)}$ time for some computable function f independent of the input size n . The rationale behind such an algorithm is that it solves the problem efficiently in applications where the parameter is small. If such an algorithm exists for a problem it is called *fixed-parameter tractable (FPT)* for p . Another option is to consider *parameterized approximation algorithms* [21, 22], which compute an α -approximation in $f(p) \cdot n^{O(1)}$ time for some parameter p .

By a result of Hochbaum and Shmoys [18], on general input graphs, a polynomial time 2-approximation algorithm exists, and this approximation factor is also best possible, unless $P=NP$. A natural parameter for k -CENTER is the number of centers k , for which however the problem is W[2]-hard [9]. In fact it is even W[2]-hard [16] to compute a $(2 - \varepsilon)$ -approximation for any $\varepsilon > 0$, and thus parametrizing by k does not help to overcome the polynomial-time inapproximability. For structural parameters such as the vertex-cover number or the feedback-vertex-set number the problem remains W[1]-hard [19], even when combining with the parameter k . For each of the more general structural parameters treewidth and cliquewidth, an *efficient parameterized approximation scheme (EPAS)* was shown to exist [19], i.e., a $(1 + \varepsilon)$ -approximation can be computed in $f(\varepsilon, w) \cdot n^{O(1)}$ time for any $\varepsilon > 0$, if w is either the treewidth or the cliquewidth, and n is the number of vertices.

Arguably however, graphs with low treewidth or cliquewidth do not model transportation networks well, since grid-like structures with large treewidth and cliquewidth can occur in road maps of big cities. As we focus on applications for k -CENTER in logistics, here we consider more natural models for transportation networks. These include *planar graphs*, *low doubling metrics* such as the Euclidean or Manhattan plane, or the more recently studied *low highway dimension graphs*. Our main result is that k -CENTER is W[1]-hard on all of these graph classes *combined*, even if adding k and the treewidth as parameters. Before introducing these graph classes, let us formally state our theorem.

► **Theorem 1.** *Even on weighted planar graphs of doubling dimension $O(1)$, the k -CENTER problem is W[1]-hard for the combined parameter (k, t, h) , where t is the treewidth and h the highway dimension of the input graph. Moreover, under ETH there is no $f(k, t, h) \cdot n^{o(t + \sqrt{k+h})}$ time algorithm³ for the same restriction on the input graphs, for any computable function f .*

A *planar graph* can be drawn in the plane without crossing edges. Such graphs constitute a realistic model for road networks, since overpasses and tunnels are relatively rare. It is known [25] that also for planar graphs no $(2 - \varepsilon)$ -approximation can be computed in polynomial time, unless $P=NP$. On the positive side, k -CENTER is FPT [9] on unweighted planar graphs for the combined parameter k and ρ . However, typically if k is small then ρ is

³ Here $o(t + \sqrt{k+h})$ means any function $g(t + \sqrt{k+h})$ such that $g(x) \in o(x)$.

large and vice versa, and thus the applications for this combined parameter are rather limited. If the parameter is only k , then an $n^{O(\sqrt{k})}$ XP-algorithm exists for planar graphs [24]. By a very recent result [20] the k -CENTER problem on weighted planar graphs admits an *efficient polynomial-time bicriteria approximation scheme*, which for any $\varepsilon > 0$ in $f(\varepsilon) \cdot n^{O(1)}$ time computes a solution that uses at most $(1 + \varepsilon)k$ centers and approximates the optimum with at most k centers within a factor of $1 + \varepsilon$. This algorithm implies an EPAS for parameter k on weighted planar graphs, since setting $\varepsilon = \min\{\varepsilon', \frac{1}{2k}\}$ forces the algorithm to compute a $(1 + \varepsilon')$ -approximation in time $f(k, \varepsilon') \cdot n^{O(1)}$ using at most $(1 + \varepsilon)k \leq k + \frac{1}{2}$ centers, i.e., at most k centers as k is an integer. This observation is complemented by our hardness result by showing that it is necessary to approximate the solution when parametrizing by k in weighted planar graphs.

► **Definition 2.** The *doubling dimension* of a metric (X, dist) is the smallest $d \in \mathbb{R}$ such that for any $r > 0$, every ball of radius $2r$ is contained in the union of at most 2^d balls of radius r . The doubling dimension of a graph is the doubling dimension of its shortest-path metric.

Since a transportation network is embedded on a large sphere (namely the earth), a reasonable model is to assume that the shortest-path metric abides to the Euclidean L_2 -norm. In cities, where blocks of buildings form a grid of streets, it is reasonable to assume that the distances are given by the Manhattan L_1 -norm. Every metric for which the distance function is given by the L_q -norm in D -dimensional space \mathbb{R}^D has doubling dimension $O(D)$. Thus a road map, which is embedded into \mathbb{R}^2 can reasonably be assumed to have constant doubling dimension. It is known [23] that k -CENTER is W[1]-hard for parameter k in two-dimensional Manhattan metrics. Also, no polynomial time $(2 - \varepsilon)$ -approximation algorithm exists for k -CENTER in two-dimensional Manhattan metrics [13], and no $(1.822 - \varepsilon)$ -approximation for two-dimensional Euclidean metrics [13]. On the positive side, Agarwal and Procopiuc [4] showed that for any L_q metric in D dimensions, the k -CENTER problem can be solved optimally in time $n^{O(k^{1-1/D})}$, and an EPAS exists for the combined parameter (ε, k, D) . We generalize the latter to any metric of doubling dimension d , as formalized by the following theorem.

► **Theorem 3.** *Given a metric of doubling dimension d , a $(1 + \varepsilon)$ -approximation for k -CENTER can be computed in $(k^k / \varepsilon^{O(kd)}) \cdot n^{O(1)}$ time.*

Theorem 1 complements this result by showing that it is necessary to approximate the cost of the solution if parametrizing by k and d .

► **Definition 4.** The *highway dimension* of a graph G is the smallest $h \in \mathbb{N}$ such that, for some universal constant $c \geq 4$, for every $r \in \mathbb{R}^+$ and every ball $B_{cr}(v)$ of radius cr , there is a set $H \subseteq B_{cr}(v)$ of *hubs* such that $|H| \leq h$ and every shortest path of length more than r lying in $B_{cr}(v)$ contains a hub of H .

The highway dimension was introduced by Abraham et al. [3] as a formalization of the empirical observation by Bast et al. [5, 6] that in a road network, starting from any point A and travelling to a sufficiently far point B along the quickest route, one is bound to pass through some member of a sparse set of “access points”, i.e., the hubs. In contrast to planar and low doubling graphs, the highway dimension has the potential to model not only road networks, but more general transportation networks such as those given by air-traffic or public transportation. This is because in such networks longer connections tend to be serviced through larger and sparser stations, which act as hubs. Abraham et al. [3] were able to prove that certain shortest-path heuristics are provably faster in low highway dimension

graphs than in general graphs. They specifically chose the constant $c = 4$ in their original definition, but later work by Feldmann et al. [14] showed that when choosing any constant $c > 4$ in the definition, the structure of the resulting graphs can be exploited to obtain quasi-polynomial time approximation schemes for problems such as TRAVELLING SALESMAN or FACILITY LOCATION. Note that increasing the constant c in Definition 4 restricts the class of graphs further. Other definitions of the highway dimension exist as well [1, 2, 3, 14] (see [15, Section 9] for a detailed discussion).

Later, Becker et al. [7] used the framework introduced by Feldmann et al. [14] to show that whenever $c > 4$ there is an EPAS for k -CENTER parameterized by ε , k , and h . Note that the highway dimension is always upper bounded by the vertex-cover number, as every edge of any non-trivial path is incident to a vertex cover. Hence the aforementioned W[1]-hardness result by Katsikarelis et al. [19] for the combined parameter k and the vertex-cover number proves that it is necessary to approximate the optimum when using k and h as the combined parameter. When parametrizing only by the highway dimension but not k , it is not even known if a *parameterized approximation scheme (PAS)* exists, i.e., an $f(\varepsilon, h) \cdot n^{g(\varepsilon)}$ time $(1 + \varepsilon)$ -approximation algorithm for some functions f, g . However, under the *Exponential Time Hypothesis (ETH)* [8], by [16] there is no algorithm with doubly exponential $2^{2^{o(\sqrt{h})}} \cdot n^{O(1)}$ runtime to compute a $(2 - \varepsilon)$ -approximation for any $\varepsilon > 0$. The same paper [16] also presents a $3/2$ -approximation for k -CENTER with runtime $2^{O(kh \log h)} \cdot n^{O(1)}$ for a more general definition of the highway dimension than the one given in Definition 4. In contrast to the result of Becker et al. [7], it is not known whether a PAS exists when combining this more general definition of h with k as a parameter. Theorem 1 complements these results by showing that even of planar graphs of constant doubling dimension, for the combined parameter (k, h) no FPT algorithm exists, unless $\text{FPT} = \text{W}[1]$. Therefore approximating the optimum is necessary, regardless of whether h is according to Definition 4 or the more general one from [16], and regardless of how restrictive Definition 4 is made by increasing the constant c .

► **Definition 5.** A *tree decomposition* of a graph $G = (V, E)$ is a tree T each of whose nodes v is labelled by a bag $X_v \subseteq V$ of vertices of G , and has the following properties:

- (a) $\bigcup_{v \in V(T)} X_v = V$,
- (b) for every edge $\{u, w\} \in E$ there is a node $v \in V(T)$ such that X_v contains both u and w ,
- (c) for every $v \in V$ the set $\{u \in V(T) \mid v \in X_u\}$ induces a connected subtree of T .

The *width* of the tree decomposition is $\max\{|X_v| - 1 \mid v \in V(T)\}$. The *treewidth* t of a graph G is the minimum width among all tree decompositions for G .

As mentioned above, arguably, bounded treewidth graphs are not a good model for transportation networks. Also it is already known that k -CENTER is W[1]-hard for this parameter, even when combining it with k [19]. We include this well-studied parameter here nonetheless, since the reduction of our hardness result in Theorem 1 implies that k -CENTER is W[1]-hard even for weighted planar graphs when *combining* any of the parameters k , h , d , and t . As noted in [15], these parameters are not bounded in terms of each other, i.e., they are incomparable. Furthermore, the doubling dimension is in fact bounded by a constant in Theorem 1. Hence, even if one were to combine all the models presented above and assume that a transportation network is planar, embeddable into some constant dimensional L_q metric, has bounded highway dimension, and even has bounded treewidth, it is unlikely that the k -CENTER problem can be solved efficiently. Thus it seems unavoidable to approximate the problem in transportation networks, when developing fast algorithms.

1.1 Related work

The k -CENTER problem is a fairly general clustering problem and therefore finds further applications in, for instance, image processing and data-compression. The above mentioned very recent efficient bicriteria approximation scheme [20] improves on a previous (non-efficient) bicriteria approximation scheme [12], which for any $\varepsilon > 0$ and weighted planar input graph computes a $(1 + \varepsilon)$ -approximation with at most $(1 + \varepsilon)k$ centers in time $n^{f(\varepsilon)}$ for some function f (note that in contrast to above, such an algorithm does not imply a PAS for parameter k). The paper by Demaine et al. [9] on the k -CENTER problem in unweighted planar graphs also considers the so-called class of *map graphs*, which is a superclass of planar graphs that is not minor-closed. They show that the problem is FPT on unweighted map graphs for the combined parameter (k, ρ) . Also for the tree-depth, k -CENTER is FPT [19]. A closely related problem to k -CENTER is the ρ -DOMINATING SET problem, in which ρ is given and the number k of centers covering a given graph with k balls of radius ρ needs to be minimized. As this generalizes the DOMINATING SET problem, no $(\ln(n) - \varepsilon)$ -approximation is possible in polynomial time [10], unless $P=NP$, and computing an $f(k)$ -approximation is $W[1]$ -hard [27] when parametrizing by k , for any computable function f .

2 The reduction

In this section we give a reduction from the GRID TILING WITH INEQUALITY (GT_{\leq}) problem, which is defined as follows. Given κ^2 non-empty sets $S_{i,j} \subseteq [n]^2$ of pairs of integers, where $i, j \in [\kappa]$, the task is to select one pair $s_{i,j} \in S_{i,j}$ for each set such that

- if $s_{i,j} = (a, b)$ and $s_{i+1,j} = (a', b')$ then $a \leq a'$, whenever $i \leq \kappa - 1$, and
- if $s_{i,j} = (a, b)$ and $s_{i,j+1} = (a', b')$ then $b \leq b'$, whenever $j \leq \kappa - 1$.

The GT_{\leq} problem is $W[1]$ -hard [8] for parameter κ , and moreover, under ETH has no $f(\kappa) \cdot n^{o(\kappa)}$ time algorithm for any computable function f .

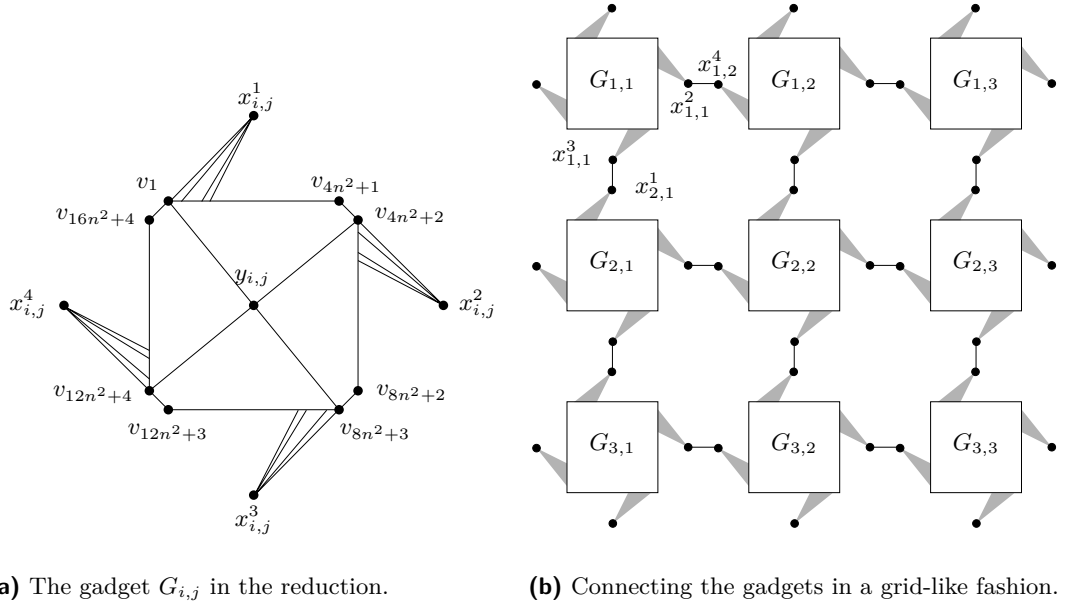
Construction

Given an instance \mathcal{I} of GT_{\leq} with κ^2 sets, we construct the following graph $G_{\mathcal{I}}$. First, for each set $S_{i,j}$, where $1 \leq i, j \leq \kappa$, we fix an arbitrary order on its elements, so that $S_{i,j} = \{s_1, \dots, s_{\sigma}\}$, where $\sigma \leq n^2$. We then construct a gadget $G_{i,j}$ for $S_{i,j}$, which contains a cycle $O_{i,j}$ of length $16n^2 + 4$ for which each edge has length 1 (see Figure 1(a)). Additionally we introduce five vertices $x_{i,j}^1, x_{i,j}^2, x_{i,j}^3, x_{i,j}^4$, and $y_{i,j}$. If $O_{i,j} = (v_1, v_2, \dots, v_{16n^2+4}, v_1)$ then we connect these five vertices to the cycle as follows. The vertex $y_{i,j}$ is adjacent to the four vertices $v_1, v_{4n^2+2}, v_{8n^2+3}$, and v_{12n^2+4} , with edges of length $2n^2 + 1$ each. For every $\tau \in [\sigma]$ and $s_{\tau} \in S_{i,j}$, if $s_{\tau} = (a, b)$ we add the four edges

- $x_{i,j}^1 v_{\tau}$ of length $\ell'_a = 2n^2 - \frac{a}{n+1}$,
- $x_{i,j}^2 v_{\tau+4n^2+1}$ of length $\ell_b = 2n^2 + \frac{b}{n+1} - 1$,
- $x_{i,j}^3 v_{\tau+8n^2+2}$ of length $\ell_a = 2n^2 + \frac{a}{n+1} - 1$, and
- $x_{i,j}^4 v_{\tau+12n^2+3}$ of length $\ell'_b = 2n^2 - \frac{b}{n+1}$.

We say that the element s_{τ} corresponds to the four vertices $v_{\tau}, v_{\tau+4n^2+1}, v_{\tau+8n^2+2}$, and $v_{\tau+12n^2+3}$. Note that s_1 (which always exists) corresponds to the four vertices adjacent to $y_{i,j}$. Note also that $2n^2 - 1 < \ell_a, \ell'_a, \ell_b, \ell'_b < 2n^2$, since $a, b \in [n]$.

The gadgets $G_{i,j}$ are now connected to each other in a grid-like fashion (see Figure 1(b)). That is, for $i \leq \kappa - 1$ we introduce a path between $x_{i,j}^3$ and $x_{i+1,j}^1$ that has $n + 2$ edges, each of length $\frac{1}{n+2}$. Analogously, for $j \leq \kappa - 1$ we add a path between $x_{i,j}^2$ and $x_{i,j+1}^4$ with $n + 2$ edges of length $\frac{1}{n+2}$ each. Note that these paths all have length 1.

(a) The gadget $G_{i,j}$ in the reduction.

(b) Connecting the gadgets in a grid-like fashion.

■ **Figure 1** The reduction.

The resulting graph $G_{\mathcal{I}}$ forms an instance of k -CENTER with $k = 5\kappa^2$. We claim that the instance \mathcal{I} of GT_{\leq} has a solution if and only if the optimum solution to k -CENTER on $G_{\mathcal{I}}$ has cost at most $2n^2$. We note at this point that the reduction would still work when removing the vertices $y_{i,j}$ and decreasing k to $4\kappa^2$. However their existence will greatly simplify analysing the doubling dimension of $G_{\mathcal{I}}$ in Section 3.

A solution to a GT_{\leq} instance implies cost at most $2n^2$ for the k -Center instance

Recall that we fixed an order of each set $S_{i,j}$, so that each element $s_{\tau} \in S_{i,j}$ corresponds to four equidistant vertices on cycle $O_{i,j}$ with distance $4n^2 + 1$ between consecutive such vertices on the cycle. If $s_{\tau} \in S_{i,j}$ is in the solution to the GT_{\leq} instance \mathcal{I} , let $C_{i,j} = \{v_{\tau}, v_{\tau+4n^2+1}, v_{\tau+8n^2+2}, v_{\tau+12n^2+3}, y_{i,j}\}$ contain the vertices of $O_{i,j}$ corresponding to s_{τ} in addition to $y_{i,j}$. The solution to the k -CENTER instance $G_{\mathcal{I}}$ is given by the union $\bigcup_{i,j \in [\kappa]} C_{i,j}$, which are exactly $5\kappa^2$ centers in total.

Let us denote the set containing the four vertices of $C_{i,j} \cap V(O_{i,j})$ by $C_{i,j}^O$, and note that each of these four vertices covers $4n^2 + 1$ vertices of $O_{i,j}$ with balls of radius $2n^2$, as each edge of $O_{i,j}$ has length 1. Since the distance between any pair of centers in $C_{i,j}^O$ is at least $4n^2 + 1$, these four sets of covered vertices are pairwise disjoint. Thus the total number of vertices covered by $C_{i,j}^O$ on $O_{i,j}$ is $16n^2 + 4$, i.e. all vertices of the cycle $O_{i,j}$ are covered. Recall that for the lengths of the edges between the vertices $x_{i,j}^1, x_{i,j}^2, x_{i,j}^3$, and $x_{i,j}^4$ and the cycle $O_{i,j}$ we have $\ell_a, \ell'_a, \ell_b, \ell'_b < 2n^2$. Hence the centers in $C_{i,j}^O$ also cover $x_{i,j}^1, x_{i,j}^2, x_{i,j}^3$, and $x_{i,j}^4$ by balls of radius $2n^2$.

Now consider a path connecting two neighbouring gadgets, e.g., let P be the path connecting $x_{i,j}^2$ and $x_{i,j+1}^4$. The center sets $C_{i,j}^O$ and $C_{i,j+1}^O$ contain vertices corresponding to the respective elements $s \in S_{i,j}$ and $s' \in S_{i,j+1}$ of the solution to the GT_{\leq} instance. This means that if $s = (a, b)$ and $s' = (a', b')$ then $b \leq b'$. Thus the closest centers of $C_{i,j}^O$ and

$C_{i,j+1}^O$ are at distance $\ell_b + 1 + \ell'_b$ from each other, as P has length 1. From $b \leq b'$ we get

$$\ell_b + 1 + \ell'_b = 2n^2 + \frac{b}{n+1} - 1 + 1 + 2n^2 - \frac{b'}{n+1} \leq 4n^2.$$

Therefore all vertices of P are covered by the balls of radius $2n^2$ around the two closest centers of $C_{i,j}^O$ and $C_{i,j+1}^O$. Analogously, we can also conclude that any path connecting some vertices $x_{i,j}^1$ and $x_{i+1,j}^3$ are covered, using the fact that if $(a, b) \in S_{i,j}$ and $(a', b') \in S_{i+1,j}$ are in the solution to the GT_{\leq} instance then $a \leq a'$.

Finally, the remaining center vertices in $\bigcup_{i,j \in [\kappa]} C_{i,j}^O \setminus C_{i,j}^O$ cover the additional vertex $y_{i,j}$ in each gadget $G_{i,j}$.

A k -Center instance with cost at most $2n^2$ implies a solution to the GT_{\leq} instance

We first prove that in any solution to the k -CENTER instance $G_{\mathcal{I}}$ of cost at most $2n^2$, each cycle $O_{i,j}$ must contain exactly four centers. Recall that $\ell_a, \ell'_a, \ell_b, \ell'_b > 2n^2 - 1$, that $y_{i,j}$ is incident to four edges of length $2n^2 + 1$ each, and that each edge of $O_{i,j}$ has length 1. Now consider the vertices $v_{4n^2+1}, v_{8n^2+2}, v_{12n^2+3}$, and v_{16n^2+4} , each of which is not connected by an edge to any vertex $x_{i,j}^h$, where $h \in [4]$, nor to $y_{i,j}$. Thus each of these four vertices must be covered by centers on the cycle $O_{i,j}$ if the radius of each ball is at most $2n^2$. Furthermore, the distance between each pair of these four vertices is at least $4n^2 + 1$, which means that any solution of cost at most $2n^2$ needs at least four centers on $O_{i,j}$ to cover these four vertices.

Each vertex $y_{i,j}$ must be contained in any solution of cost at most $2n^2$, since the distance from $y_{i,j}$ to any other vertex is more than $2n^2$. This already uses κ^2 of the available $5\kappa^2$ centers. Since there are κ^2 cycles and only $4\kappa^2$ remaining available centers, we proved that each cycle $O_{i,j}$ contains exactly four centers, and no other centers exist in the graph $G_{\mathcal{I}}$. Let $C_{i,j}^O$ be the set of four centers contained in $O_{i,j}$. As each center of $C_{i,j}^O$ covers at most $4n^2 + 1$ vertices of $O_{i,j}$ by balls of radius at most $2n^2$, to cover all $16n^2 + 4$ vertices of $O_{i,j}$ these four centers must be equidistant with distance exactly $4n^2 + 1$ between consecutive centers on $O_{i,j}$. Furthermore, since $\ell_a, \ell'_a, \ell_b, \ell'_b > 2n^2 - 1$ and each edge of $O_{i,j}$ has length 1, to cover $x_{i,j}^h$ for any $h \in [4]$ some center of $C_{i,j}^O$ must lie on a vertex of $O_{i,j}$ adjacent to $x_{i,j}^h$. This means that the four centers of $C_{i,j}^O$ are exactly those vertices $v_{\tau+(h-1)(4n^2+1)}$ corresponding to element s_{τ} of $S_{i,j}$.

It remains to show that the elements corresponding to the centers in $\bigcup_{i,j \in [\kappa]} C_{i,j}^O$ form a solution to the GT_{\leq} instance \mathcal{I} . For this, consider two neighbouring gadgets $G_{i,j}$ and $G_{i,j+1}$, and let $(a, b) \in S_{i,j}$ and $(a', b') \in S_{i,j+1}$ be the respective elements corresponding to the center sets $C_{i,j}^O$ and $C_{i,j+1}^O$. Note that for any $\hat{b} \in [n]$ we have $\ell_b \leq \ell_{\hat{b}} + 1$ and $\ell'_b \leq \ell'_{\hat{b}} + 1$. Since every edge of the cycles $O_{i,j}$ and $O_{i,j+1}$ has length 1, this means that the distance from the closest centers $v \in C_{i,j}^O$ and $v' \in C_{i,j+1}^O$ to $x_{i,j}^2$ and $x_{i,j+1}^4$, respectively, is determined by the edges of length ℓ_b and ℓ'_b incident to v and v' , respectively. In particular, the distance between v and v' is $\ell_b + 1 + \ell'_b$, as the path P connecting $x_{i,j}^2$ and $x_{i,j+1}^4$ has length 1. Assume now that $b > b'$, which means that $b \geq b' + 1$ since b and b' are integer. Hence this distance is

$$\ell_b + 1 + \ell'_b = 2n^2 + \frac{b}{n+1} - 1 + 1 + 2n^2 - \frac{b'}{n+1} \geq 4n^2 + \frac{1}{n+1}.$$

As the centers v and v' only cover vertices at distance at most $2n^2$ each, while the edges of the path P have length $\frac{1}{n+2} < \frac{1}{n+1}$, there must be some vertex of P that is not covered by the center set. However this contradicts the fact that the centers form a feasible solution with cost at most $2n^2$, and so $b \leq b'$.

An analogous argument can be made for neighbouring gadgets $G_{i,j}$ and $G_{i+1,j}$, so that $a \leq a'$ for the elements $(a,b) \in S_{i,j}$ and $(a',b') \in S_{i+1,j}$ corresponding to the centers in $C_{i,j}^O$ and $C_{i+1,j}^O$, respectively. Thus a solution to $G_{\mathcal{I}}$ of cost at most $2n^2$ implies a solution to \mathcal{I} .

3 Properties of the constructed graph

The reduction of Section 2 proves that the k -CENTER problem is W[1]-hard for parameter k , since the reduction can be done in polynomial time and k is function of κ . Since this function is quadratic, we can also conclude that, under ETH, there is no $f(k) \cdot n^{o(\sqrt{k})}$ algorithm for k -CENTER. We will now show that the reduction has various additional properties from which we will be able to conclude Theorem 1. First of all we prove that any constructed graph $G_{\mathcal{I}}$ for an instance \mathcal{I} of GT_{\leq} is planar and has bounded doubling dimension.

► **Lemma 6.** *The graph $G_{\mathcal{I}}$ is planar and has doubling dimension at most $\log_2(36) \approx 5.17$ for $n \geq 3$.*

Proof. It is obvious from the construction of the graph $G_{\mathcal{I}}$ that it is planar. To bound its doubling dimension, consider the shortest-path metric on the vertex set $Y = \{y_{i,j} \in V(G_{\mathcal{I}}) \mid i, j \in [\kappa]\}$ given by the distances between these vertices in $G_{\mathcal{I}}$. As these vertices are arranged in a grid-like fashion, this shortest-path metric on Y approximates the L_1 -metric. We consider a set of index pairs, for which the corresponding vertices in Y roughly resemble a ball in the shortest-path metric on Y . That is, consider the set of index pairs $A_{i,j}(a) = \{(i',j') \in [\kappa]^2 \mid |i - i'| + |j - j'| \leq a\}$, and let $V_{i,j}(a) \subseteq V(G_{\mathcal{I}})$ contain all vertices of gadgets $G_{i',j'}$ such that $(i',j') \in A_{i,j}(a)$ in addition to the vertices of paths of length 1 connecting these gadgets. We would like to bound the diameter of the graph induced by $V_{i,j}(a)$ in $G_{\mathcal{I}}$, for which we need the following claim, which we will reuse later.

► **Claim 7.** *For any gadget $G_{i,j}$ and $h, h' \in [4]$ with $h \neq h'$, the distance between $x_{i,j}^h$ and $x_{i,j}^{h'}$ lies between $7n^2 - 1$ and $8n^2 + 2$.*

Proof. The distance between $x_{i,j}^h$ and $x_{i,j}^{h'}$ is less than $2(2n^2 + 2n^2 + 1) = 8n^2 + 2$, via the path passing through $y_{i,j}$ and the two vertices of $O_{i,j}$ adjacent to $y_{i,j}$, $x_{i,j}^h$, and $x_{i,j}^{h'}$. Note that the shortest path between $x_{i,j}^h$ and $x_{i,j}^{h'}$ inside the gadget $G_{i,j}$ does not necessarily pass through $y_{i,j}$, but may pass along the cycle $O_{i,j}$ instead. This is because the set $S_{i,j}$ of the GT_{\leq} instance may contain up to n^2 elements, which would imply a direct edge from $x_{i,j}^h$ to $v_{n^2+(h-1)(4n^2+1)}$ on $O_{i,j}$. Thus we can give a lower bound of $2(2n^2 - 1) + 3n^2 + 1 = 7n^2 - 1$ for the distance between $x_{i,j}^h$ and $x_{i,j}^{h'}$ inside of $G_{i,j}$. This is also the shortest path between these vertices in $G_{\mathcal{I}}$, since any other path needs to pass through at least three gadgets. ◀

By Claim 7, the diameter of $V_{i,j}(a)$ is at most $(8n^2 + 3)(2a + 1)$ as one needs to traverse at most $2a + 1$ gadgets $G_{i',j'}$ with $(i',j') \in A_{i,j}(a)$ and the paths of length 1 connecting them, in order to reach any vertex of $V_{i,j}(a)$ from any other. Assuming $|A_{i,j}(a)| = (2a + 1)^2$, i.e., the set contains the maximum number of index pairs, the diameter of $V_{i,j}(a)$ is at least $7n^2(2a + 1) - 1$, since any path between two points of $V_{i,j}(a)$ at maximum distance must pass through at least $2a + 1$ gadgets $G_{i',j'}$ with $(i',j') \in A_{i,j}(a)$ and the $2a$ paths of length 1 connecting them.

Consider a ball $B_v(2r)$ around a vertex v of radius $2r$ in $G_{\mathcal{I}}$, and let $y_{i,j}$ be the closest vertex of Y to v . The distance between $y_{i,j}$ and v is at most $2(2n^2 + 1) = 4n^2 + 2$, whether v lies on $O_{i,j}$ or on one of the paths of length 1 connecting $G_{i,j}$ with an adjacent gadget. Hence the ball $B_v(2r)$ is contained in a ball of radius $4n^2 + 2 + 2r$ around $y_{i,j}$. The latter ball

is in turn contained in the $V_{i,j}(a)$ set centered at $y_{i,j}$ if its diameter is at most the diameter of $V_{i,j}(a)$, i.e., $2(4n^2 + 2 + 2r) \leq 7n^2(2a + 1) - 1$. This for instance is true if $2a + 1 = \frac{r}{n^2}$, $r \geq 4n^2 + 2$, and $n \geq 1$. From the upper bound on the diameter of a set $V_{i',j'}(a')$, we also know that $V_{i',j'}(a')$ is contained in a ball of radius r around $y_{i',j'}$ if $(8n^2 + 3)(2a' + 1) \leq 2r$, which is true if $2a' + 1 = \frac{2r}{8n^2 + 3}$. However we also want $V_{i',j'}(a')$ to be non-empty, i.e., $a' \geq 0$, which by the latter equality means that $r \geq 4n^2 + 3/2$. We may cover all vertices of $A_{i,j}(a)$ with $\lceil \frac{2a+1}{2a'+1} \rceil^2$ sets $A_{i',j'}(a')$, since in Y these sets correspond to “squares rotated by 45 degrees of diameter $2a + 1$ and $2a' + 1$, respectively”. Thus we can cover $V_{i,j}(a)$ with $\lceil \frac{2a+1}{2a'+1} \rceil^2$ sets $V_{i',j'}(a')$, i.e., if $r \geq 4n^2 + 2$ and $n \geq 2$ we can cover a ball of radius $2r$ in $G_{\mathcal{I}}$ with $\lceil \frac{2a+1}{2a'+1} \rceil^2 = \lceil 4 + \frac{3}{2n^2} \rceil^2 \leq 25$ balls of radius r .

$G_{i,j}$ gadgets $y_{i,j}$. is $r < 4 + 1)^2 = 81$ cover a

If $r < 4n^2 + 2$, a ball $B_v(2r)$ has radius less than $8n^2 + 4$. Consider the case when v lies in a gadget $G_{i,j}$ of $G_{\mathcal{I}}$. The distance from $G_{i,j}$ to any cycle $O_{i',j'}$ for which $|i - i'| + |j - j'| \geq 2$ is at least $7n^2 - 1 + 2n^2 - 1 \geq 8n^2 + 4$, as $n \geq 3$: to reach $O_{i',j'}$ a path from $G_{i,j}$ first needs to traverse a neighbouring gadget of $G_{i,j}$, which we know has diameter at least $7n^2 - 1$ by Claim 7, and the vertex $x_{i',j'}^h$ has distance more than $2n^2 - 1$ from $O_{i',j'}$. Thus $B_v(2r)$ contains at most the four neighbouring gadgets of $G_{i,j}$ and the paths of length 1 connected to these. On each of the five cycles $O_{i',j'}$ that intersect $B_v(2r)$, at most 3 balls of radius r are needed to cover all vertices in the intersection of $O_{i',j'}$ and $B_v(2r)$: as the edges of $O_{i',j'}$ all have length 1 we may choose 3 vertices equidistantly at every $\lfloor 2r \rfloor$ -th vertex on the part of $O_{i',j'}$ in $B_v(2r)$. As long as $r \geq 1$, any path of length 1 that intersects $B_v(2r)$ can be covered by one ball of radius r . Any vertex $y_{i',j'}$ that lies in $B_v(2r)$ can also be covered by one ball of radius r . As $B_v(2r)$ intersects at most 5 cycles $O_{i',j'}$, as well as at most 5 vertices $y_{i',j'}$, and at most 16 paths of length 1, these amount to at most 36 balls of radius r .

If v does not lie in any gadget $G_{i,j}$, then it lies on some path of length 1 connecting two gadgets. Given that $r < 4n^2 + 2$, in this case the ball $B_v(2r)$ intersects at most 2 cycles $O_{i,j}$ and vertices $y_{i,j}$, and at most 7 paths of length 1, since by Claim 7 the diameter of a gadget is at least $7n^2 - 1 \geq r$ if $n \geq 1$. Thus in this case at most 17 balls of radius r suffice.

Finally, if $r < 1$, then a ball $B_v(2r)$ contains only a subpath of some cycle $O_{i,j}$, a subpath of a path of length 1 connecting two gadgets, or a single vertex $y_{i,j}$, since any edge connecting a cycle $O_{i,j}$ to $y_{i,j}$ or some $x_{i,j}^h$ has length more than $2n^2 \geq 1$ if $n \geq 1$. In this case at most 3 balls of radius r suffice to cover all vertices of $B_v(2r)$. ◀

We next show that we can bound the parameters t and h , i.e. the treewidth and highway dimension of $G_{\mathcal{I}}$, as a function of the parameter $k = \Theta(\kappa^2)$. Note that the following lemma bounds the highway dimension in terms of k , no matter how restrictive we make Definition 4 by increasing the constant c .

► **Lemma 8.** *For any constant c of Definition 4, the graph $G_{\mathcal{I}}$ has highway dimension at most $O(\kappa^2)$.*

Proof. For any scale $r \in \mathbb{R}^+$ and universal constant $c \geq 4$ we will define a hub set $H_r \subseteq V$ hitting all shortest paths of length more than r in $G_{\mathcal{I}}$, such that $|H_r \cap B_{cr}(v)| = O(\kappa^2)$ for any ball $B_{cr}(v)$ of radius cr in $G_{\mathcal{I}}$. This bounds the highway dimension to $O(\kappa^2)$ according to Definition 4.

Let $X = \{y_{i,j}, x_{i,j}^h \mid h \in \{1, 2, 3, 4\} \wedge i, j \in [\kappa]\}$, i.e. it contains all vertices connecting gadgets $G_{i,j}$ to each other in addition to the vertices $y_{i,j}$. If $r > 8n^2 + 2$ then $H_r = X$. Any shortest path containing only vertices of a cycle $O_{i,j}$ has length at most $8n^2 + 2$, since the cycle has length $16n^2 + 4$. Any (shortest) path that is a subpath of a path connecting two gadgets has length at most 1. Hence any shortest path of length more than $8n^2 + 2$ must

contain some vertex of X . The total size of X is $5\kappa^2$, and so any ball, no matter its radius, also contains at most this many hubs of H_r .

If $1 \leq r \leq 8n^2 + 2$ then any path of length more than r but not containing any vertex of X must lie on some cycle $O_{i,j} = (v_1, v_2, \dots, v_{16n^2+4}, v_1)$. We define the set $W_{i,j} = \{v_{1+\lambda[r]} \in V(O_{i,j}) \mid \lambda \in \mathbb{N}_0\}$, i.e. it contains every r -th vertex on the cycle after rounding down. This means that every path on $O_{i,j}$ of length more than r contains a vertex of $W_{i,j}$. Thus for these values of r we set $H_r = X \cup \bigcup_{i,j \in [\kappa]} W_{i,j}$. Any ball $B_{cr}(v)$ of radius cr contains $O(c)$ hubs of any $W_{i,j}$. By Claim 7, the distance between any pair of the four vertices $x_{i,j}^h$, where $h \in \{1, 2, 3, 4\}$, that connect a gadget $G_{i,j}$ to other gadgets, is more than $7n^2 - 1$. This means that $B_{cr}(v)$ can only intersect $O(c^2)$ gadgets, since $cr \leq c(8n^2 + 2)$ and the gadgets are connected in a grid-like fashion. Hence the ball $B_{cr}(v)$ only contains $O(c)$ hubs for each of the $O(c^2)$ sets $W_{i,j}$ for which $B_{cr}(v)$ intersect the respective gadget $G_{i,j}$. At the same time each gadget contains only 5 vertices of X . Thus if c is a constant, then the number of hubs of H_r in $B_{cr}(v)$ is constant.

If $r < 1$, a path of length more than r may be a subpath of a path connecting two gadgets. Let $P_{i,j}$ be the path connecting $x_{i,j}^2$ and $x_{i,j+1}^4$ for $j \leq \kappa - 1$, and let $P'_{i,j}$ be the path connecting $x_{i,j}^3$ and $x_{i+1,j}^1$ for $i \leq \kappa - 1$. Recall that each of these paths consists of $n+2$ edges of length $\frac{1}{n+2}$ each. If $P_{i,j} = (u_0, u_1, \dots, u_{n+2})$, we define the set $U_{i,j} = \{u_{\lambda[r](n+2)} \in V(P_{i,j}) \mid \lambda \in \mathbb{N}_0\}$, and if $P'_{i,j} = (u_0, u_1, \dots, u_{n+2})$, we define the set $U'_{i,j} = \{u_{\lambda[r](n+2)} \in V(P'_{i,j}) \mid \lambda \in \mathbb{N}_0\}$, i.e. these sets contain vertices of consecutive distance r on the respective paths, after rounding down. Now let $H_r = \bigcup_{i,j \in [\kappa]} V(G_{i,j}) \cup U_{i,j} \cup U'_{i,j}$, so that every path of length more than r contains a hub of H_r . Any ball $B_{cr}(v)$ of radius $cr < c$ intersects only $O(c^2)$ gadgets $G_{i,j}$, as observed above. As the edges of a cycle $O_{i,j}$ have length 1, the ball B contains only $O(c)$ vertices of $O_{i,j}$. Thus $B_{cr}(v)$ contains $O(c)$ hubs of $V(G_{i,j}) \cup U_{i,j} \cup U'_{i,j}$ for each of the $O(c^2)$ gadgets $G_{i,j}$ it intersects. For constant c , this proves the claim. \blacktriangleleft

To bound the treewidth of $G_{\mathcal{I}}$ we use so-called *cops and robber games*. Given a graph G and $\tau \in \mathbb{N}$, a *state* of the τ *cops and robber game* on G is a pair (K, p) where $K \subseteq V$ with $|K| \leq \tau$, and $p \in V \setminus K$. The set K encodes the positions of the τ cops, while p is the position of the robber. The game proceeds in rounds, where each round $z \in \mathbb{N}_0$ is associated with a state (K_z, p_z) . Initially, in round 0 the cops first choose positions K_0 and then the robber chooses a position $p_0 \in V \setminus K_0$. In each round $z \geq 1$, first the cops choose a new position K_z , after which the robber can choose a position $p_z \in V \setminus K_z$, such that p_z and p_{z-1} lie in the same connected component of $G - (K_z \cap K_{z-1})$. The cops win the game if after a finite number of rounds, the robber has no position to choose, i.e., the robber is caught. By [26], a graph G has treewidth t if and only if $t + 1$ cops can win in G .

► **Lemma 9.** *The graph $G_{\mathcal{I}}$ has treewidth at most $\kappa + O(1)$.*

Proof. We prove that $\kappa + O(1)$ cops can win the cops and robber game on $G_{\mathcal{I}}$. For each $i, j \in [\kappa]$ we define the sets $X_{i,j}^2 = \{x_{i',j}^2 \mid i' \in [i]\}$ and $X_{i,j}^4 = \{x_{i',j}^4 \mid i' \in [\kappa] \setminus [i-1]\}$ and let $K_{i,j} = \{y_{i,j}, x_{i,j}^1, x_{i,j}^3\} \cup X_{i,j}^2 \cup X_{i,j}^4$ be a position for the cops. The initial position of the cops is $K_{1,1}$, and they will sweep the graph $G_{\mathcal{I}}$ “from left to right” with increasing index j . More precisely we will describe a strategy, which uses a finite number of intermediate rounds to go from position $K_{i,j}$ to position $K_{i+1,j}$ for each $i \leq \kappa - 1$, and from position $K_{\kappa,j}$ to position $K_{1,j+1}$ for each $j \leq \kappa - 1$. After reaching position $K_{\kappa,\kappa}$ the robber will be caught.

Consider a position $K_{i,j}$ and note that the three connected components left after removing all vertices of $K_{i,j}$ from $G_{\mathcal{I}}$ are (a) the cycle $O_{i,j}$, (b) the subgraph $L_{i,j}$ “left of” $G_{\mathcal{I}}$ induced by all gadgets $G_{i',j'}$ and paths $P_{i',j'}, P'_{i',j'}$ for which $j' \leq j - 1$ and $i' \leq \kappa$, but also the gadgets $G_{i',j'}$ and paths $P'_{i',j'}$ for which $j' = j$ and $i' \leq i - 1$, and finally (c) the subgraph

$R_{i,j}$ “right of” $G_{\mathcal{I}}$ induced by all gadgets $G_{i',j'}$ and paths $P_{i',j'}, P'_{i',j'}$ for which either $j' = j$ and $i' \geq i+1$, or $j' \geq j+1$ and $i' \leq \kappa$, but also the paths $P_{i',j}$ where $i' \leq i$ and the path $P'_{i,j}$. The robber’s position p has to be in one of these subgraphs, and we will describe a strategy of the cops, which guarantees that p is a vertex of $R_{i,j}$ if the robber has not been caught yet. As $R_{\kappa,\kappa}$ is empty, this means that the cops are able to catch the robber eventually.

In the initial position $K_{1,1}$ the position p cannot be in $L_{1,1}$ as this graph is empty. We now show how in any position $K_{i,j}$ the cycle $O_{i,j}$ can be traversed by the cops using only three additional cops to those in $K_{i,j}$, to catch the robber in case his position p is on $O_{i,j}$. If $O_{i,j} = (v_1, v_2, \dots, v_{16n^2+4}, v_1)$ we define a sequence of positions $K^h = K_{i,j} \cup \{v_1, v_h, v_{h+1}\}$ where $h \in [16n^2 + 3]$. Note that $v_1, v_{h+1} \in K^h \cap K^{h+1}$ so that $O_{i,j} - (K^h \cap K^{h+1})$ consists of the two paths $Q_1^h = (v_2, v_3, \dots, v_h)$ and $Q_2^h = (v_{h+2}, v_{h+3}, \dots, v_{16n^2+4})$. The former path Q_1^h is empty for $h = 1$ and thus the robber cannot move to a position on Q_1^h throughout the whole sequence in any step h . In the last step $h = 16n^2 + 3$ however, $Q_2^{16n^2+3}$ is empty, and so the robber would not have any position to choose in this step if p was on Q_2^h on any previous step h . Thus we may assume that the robber’s position lies in $R_{i,j}$ by induction.

If the cops are in position $K_{i,j}$ for some $i \leq \kappa - 1$, then they can move to position $K_{i+1,j}$ as follows. The cops first move to position $K_{i,j} \cup K_{i+1,j}$, which contains $\kappa + 6$ vertices. By induction, the position p of the robber is now in $R_{i,j} - K_{i+1,j}$, i.e., p lies either in $R_{i+1,j}$, $O_{i+1,j}$, or $P'_{i,j}$. As described above, the cops can catch the robber if p is on the cycle $O_{i,j}$. A similar strategy can be used to traverse the path $P'_{i,j}$: if $P'_{i,j} = (u_0, u_1, \dots, u_{n+2})$ then we define a sequence of positions $K^h = K_{i,j} \cup K_{i+1,j} \cup \{u_{h-1}, u_h\}$ for $h \in [n+1]$, so that $P'_{i,j} - (K^h \cap K^{h+1})$ consists of the two paths $Q_1^h = (u_1, u_2, \dots, u_{h-1})$ and $Q_2^h = (u_{h+1}, u_{h+2}, \dots, u_{n+1})$, since $u_0 = x_{i,j}^3 \in K_{i,j}$ and $u_{n+2} = x_{i+1,j}^1 \in K_{i+1,j}$. As Q_1^1 is empty and $u_h \in K^h \cap K^{h+1}$, the robber’s position cannot be on Q_1^h in any subsequent step h , and as Q_2^{n+1} is empty, the robber is caught if p lies on path $P'_{i,j}$. Thus the only possibility left for the robber not to be caught is to be in $R_{i+1,j}$. In this case the cops can switch to position $K_{i+1,j}$ and continue their chase.

Finally the cops also need a strategy to move from position $K_{\kappa,j}$ to position $K_{1,j+1}$ for each $j \leq \kappa - 1$. For this we define the intermediate positions $K'_{i,j} = X_{i,j}^2 \cup X_{i,j}^4$, which contain $\kappa + 1$ cops each. Note that $R_{\kappa,j}$ is a connected component of $G_{\mathcal{I}} - (K_{\kappa,j} \cap K'_{\kappa,j})$. Thus the cops can safely switch from position $K_{\kappa,j}$ to $K'_{\kappa,j}$, as the robber is in $R_{\kappa,j}$ by induction. For any $i \in [\kappa]$, removing the vertices of $K'_{i,j}$ from $G_{\mathcal{I}}$ leaves three connected components of which one is $P_{i,j}$ without the endpoints, and one is a component $R'_{i,j}$, which is $R_{i,j}$ without the paths $P_{i',j}$ where $i' \geq i$. By induction the robber’s position p is in one of these two components. If p is on $P_{i,j}$, the above strategy for paths shows how to catch the robber. Thus the only possibility left is that p lies in $R'_{i,j}$. If $i \geq 2$, the cops can switch to position $K'_{i-1,j}$, after which p again lies either on $P_{i-1,j}$ or in $R'_{i-1,j}$. For $i = 1$, note that $L_{1,j+1}$ is a connected component of $G_{\mathcal{I}} - (K_{1,j+1} \cap K'_{1,j})$, which by induction does not contain the robber. Thus when switching from position $K'_{1,j}$ to $K_{1,j+1}$ the robber must either be on $O_{1,j+1}$ or $R_{1,j+1}$. We know how to catch the robber if he happens to be on $O_{1,j+1}$ using only three additional cops, and so the only case left is that p is in $R_{1,j+1}$.

Each of the used positions for the cops has at most $\kappa + O(1)$ vertices, and thus the treewidth of $G_{\mathcal{I}}$ is bounded by the same term. \blacktriangleleft

The reduction given in Section 2 together with Lemmas 6, 8 and 9 imply Theorem 1, since the GT_{\leq} problem is $\text{W}[1]$ -hard [8] for parameter κ , and we may assume w.l.o.g. that $n \geq 3$. Moreover $\kappa = \Theta(k)$ and, under ETH, GT_{\leq} has no $f(\kappa) \cdot n^{o(\kappa)}$ time algorithm [8] for any computable function f .

4 An algorithm for low doubling metrics

In this section we give a simple algorithm that generalizes one from [4], which for D -dimensional L_q metrics compute a $(1 + \varepsilon)$ -approximation in time $f(\varepsilon, k, D) \cdot n^{O(1)}$. In particular, any such metric has doubling dimension $O(D)$. Here we assume that the input metric has doubling dimension d . A fundamental observation about metrics of bounded doubling dimension is the following, which can be proved by a simple recursive application of Definition 2. Here the aspect ratio is the diameter of X divided by the minimum distance of the metric.

► **Lemma 10** ([17]). *Let (X, dist) be a metric with doubling dimension d and $Y \subseteq X$ be a subset with aspect ratio α . Then $|Y| \leq 2^{d \lceil \log_2 \alpha \rceil}$.*

To compute a $(1 + \varepsilon)$ -approximation to k -CENTER given a graph G with vertex set V , we first compute its shortest-path metric (V, dist) . We then compute a *net* of this metric, which is defined as follows.

► **Definition 11.** For a metric (X, dist) , a subset $Y \subseteq X$ is called a δ -cover if for every $u \in X$ there is a $v \in Y$ such that $\text{dist}(u, v) \leq \delta$. A δ -net is a δ -cover with the additional property that $\text{dist}(u, v) > \delta$ for all vertices $u, v \in Y$.

Note that a δ -net can be computed greedily in polynomial time. The first step of our algorithm is to guess the optimum cost ρ by trying each of the $\binom{n}{2}$ possible values. For each guess we compute an $\frac{\varepsilon\rho}{2}$ -net $Y \subseteq X$. We know that the metric (V, dist) can be covered by k balls of diameter 2ρ each, which means that the aspect ratio of Y inside of each ball is at most $4/\varepsilon$. Thus by Lemma 10, each ball contains $1/\varepsilon^{O(d)}$ vertices of Y , and so $|Y| \leq k/\varepsilon^{O(d)}$.

An optimum k -CENTER solution $C \subseteq Y$ for (Y, dist) can be computed by brute force in $\binom{|Y|}{k} = k^k/\varepsilon^{O(kd)}$ steps. Since every center of the optimum solution $C^* \subseteq V$ of the input graph has a net point of Y at distance at most $\frac{\varepsilon\rho}{2}$, there exists a k -CENTER solution in Y of cost at most $(1 + \varepsilon/2)\rho$. The computed center set $C \subseteq Y$ thus also has cost at most $\frac{\varepsilon\rho}{2}$. Therefore C covers all of V with balls of radius $(1 + \varepsilon)\rho$, since every vertex of V is at distance $\frac{\varepsilon\rho}{2}$ from some vertex of Y . Thus C is a $(1 + \varepsilon)$ -approximation of the input graph. Considering the guessed values of ρ in increasing order, outputting the first computed solution with at most k centers gives the algorithm of Theorem 3.

References

- 1 I. Abraham, D. Delling, A. Fiat, A. V. Goldberg, and R. F. Werneck. Highway dimension and provably efficient shortest path algorithms. *Journal of the ACM*, 63(5):41, 2016.
- 2 I. Abraham, D. Delling, A. Fiat, A.V. Goldberg, and R.F. Werneck. VC-dimension and shortest path algorithms. In *ICALP*, pages 690–699, 2011.
- 3 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.
- 4 P. K. Agarwal and C. M. Procopiuc. Exact and approximation algorithms for clustering. *Algorithmica*, 33(2):201–226, 2002.
- 5 H. Bast, S. Funke, and D. Matijevic. Ultrafast shortest-path queries via transit nodes. *9th DIMACS Implementation Challenge*, 74:175–192, 2009.
- 6 H. Bast, S. Funke, D. Matijevic, P. Sanders, and D. Schultes. In transit to constant time shortest-path queries in road networks. In *ALENEX*, pages 46–59, 2007.
- 7 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. *ArXiv e-prints*, arXiv:1707.08270 [cs.DS], 2017.

- 8 M. Cygan, F. V. Fomin, L. Kowalik, D. Lokshtanov, D. Marx, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. *Parameterized Algorithms*. Springer, 2015.
- 9 E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Fixed-parameter algorithms for (k, r) -center in planar graphs and map graphs. *Transactions on Algorithms*, 1(1):33–47, 2005.
- 10 Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *STOC*. ACM Press, 2014.
- 11 R. G. Downey and M. R. Fellows. *Fundamentals of parameterized complexity*. Springer, 2013.
- 12 David Eisenstat, Philip N Klein, and Claire Mathieu. Approximating k -center in planar graphs. In *SODA*, pages 617–627, 2014.
- 13 T. Feder and D. Greene. Optimal algorithms for approximate clustering. In *STOC*, pages 434–444, 1988.
- 14 A. E. Feldmann, W. S. Fung, J. Könemann, and I. Post. A $(1+\varepsilon)$ -embedding of low highway dimension graphs into bounded treewidth graphs. In *ICALP*, pages 469–480, 2015.
- 15 A. E. Feldmann, W. S. Fung, J. Könemann, and I. Post. A $(1+\varepsilon)$ -embedding of low highway dimension graphs into bounded treewidth graphs. *ArXiv preprint arXiv:1502.04588*, 2015.
- 16 Andreas Emil Feldmann. Fixed parameter approximations for k -center problems in low highway dimension graphs. In *ICALP*, pages 588–600. Springer Berlin Heidelberg, 2015.
- 17 A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *FOCS*, pages 534–543, 2003.
- 18 D. S. Hochbaum and D. B. Shmoys. A unified approach to approximation algorithms for bottleneck problems. *Journal of the ACM*, 33(3):533–550, 1986.
- 19 Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Structural parameters, tight bounds, and approximation for (k, r) -center. In *ISAAC*, pages 50:1–50:13, 2017.
- 20 P. Klein. Personal communication, 2017.
- 21 Daniel Lokshtanov, Fahad Panolan, M. S. Ramanujan, and Saket Saurabh. Lossy kernelization. In *STOC*, pages 224–237. ACM Press, 2017.
- 22 D. Marx. Parameterized complexity and approximation algorithms. *The Computer Journal*, 51(1):60–78, 2008.
- 23 Dániel Marx. Efficient approximation schemes for geometric problems? In *European Symposium on Algorithms*, pages 448–459. Springer, 2005.
- 24 Dániel Marx and Michał Pilipczuk. Optimal parameterized algorithms for planar facility location problems using Voronoi diagrams. In *ESA*, pages 865–877. Springer, 2015.
- 25 J. Plesník. On the computational complexity of centers locating in a graph. *Aplikace matematiky*, 25(6):445–452, 1980.
- 26 Paul D Seymour and Robin Thomas. Graph searching and a min-max theorem for tree-width. *Journal of Combinatorial Theory, Series B*, 58(1):22–33, 1993.
- 27 Karthik Srikanta, Bundit Laekhanukit, and Pasin Manurangsi. On the parameterized complexity of approximating dominating set. *arXiv preprint*, abs/1711.11029, 2017.
- 28 V. V. Vazirani. *Approximation Algorithms*. Springer-Verlag New York, Inc., 2001.
- 29 David P Williamson and David B Shmoys. *The design of approximation algorithms*. Cambridge university press, 2011.