Algorithmica (2021) 83:1352–1370
https://doi.org/10.1007/s00453-020-00785-5

Travelling on Graphs with Small Highway Dimension

Yann Disser[1] · Andreas Emil Feldmann[2] · Max Klimm[3] · Jochen Könemann[4]

Received: 3 October 2019 / Accepted: 16 November 2020 / Published online: 12 February 2021

Abstract

We study the Travelling Salesperson (TSP) and the Steiner Tree problem (STP) in graphs of low highway dimension. This graph parameter roughly measures how many central nodes are visited by all shortest paths of a certain length. It has been shown that transportation networks, on which TSP and STP naturally occur for various applications in logistics, typically have a small highway dimension. While it was previously shown that these problems admit a quasi-polynomial time approximation scheme on graphs of constant highway dimension, we demonstrate that a significant improvement is possible in the special case when the highway dimension is 1. Specifically, we present a fully-polynomial time approximation scheme (FPTAS). We also prove that both TSP and STP are weakly NP-hard for these restricted graphs.

Keywords Travelling Salesperson · Steiner Tree · Highway dimension · Approximation scheme · NP-Hardness

✉ Andreas Emil Feldmann
feldmann.a.e@gmail.com

Yann Disser
disser@mathematik.tu-darmstadt.de

Max Klimm
klimm@tu-berlin.de

Jochen Könemann
jochen@uwaterloo.ca

[1] Department of Mathematics, TU Darmstadt, Doliviostr. 15, 64293 Darmstadt, Germany

[2] Department of Applied Mathematics, Charles University in Prague, Malostranske nam. 25, 118 00 Praha 1, Czechia

[3] Institute of Mathematics, Technische Universität Berlin, Str. des 17. Juni 136, 10623 Berlin, Germany

[4] Faculty of Mathematics, University of Waterloo, 200 University Ave West, Waterloo, ON N2L 3G1, Canada

1 Introduction

Two fundamental optimization problems already included in Karp's initial list of 21 NP-complete problems [38] are the Travelling Salesperson problem (TSP) and the Steiner Tree problem (STP). Given an undirected graph $G = (V, E)$ with non-negative edge weights $w : E \to \mathbb{R}^+$, the TSP asks to find the shortest closed walk in G visiting all nodes of V. Besides its fundamental role in computational complexity and combinatorial optimization, this problem has a variety of applications ranging from circuit manufacturing [33, 46] and scientific imaging [13] to vehicle routing problems [45] in transportation networks. For the STP, a subset $R \subseteq V$ of nodes is marked as *terminals*. The task is to find a weight-minimal connected subgraph of G containing the terminals. It has plenty of fundamental applications in network design including telecommunication networks [47], computer vision [19], circuit design [34], and computational biology [21, 48], but also lies at the heart of line planning in public transportation [16].

Both TSP and STP are APX-hard in general [6, 12, 20, 39, 44, 51] implying that, unless P = NP, none of these problems admit a *polynomial-time approximation scheme (PTAS)*, i.e., an algorithm that computes a $(1 + \varepsilon)$-approximation in polynomial time for any given constant $\varepsilon > 0$. On the other hand, for restricted inputs PTASs do exist, e.g., for planar graphs [5, 17, 32, 41], Euclidean and Manhattan metrics [5, 49], and more generally low doubling[1] metrics [7].

We study another class of graphs captured by the notion of *highway dimension* proposed by Abraham et al. [3]. This graph parameter models transportation networks and is thus of particular importance in terms of applications for both TSP and STP. On a high level, the highway dimension is based on the empirical observation of Bast et al. [8, 9] that travelling from a point in a network to a sufficiently distant point on a shortest path always passes through a sparse set of "hubs". The following formal definition is taken from [25] and follows the lines of Abraham et al. [3].[2] Here the *distance* between two vertices is the length of the shortest path between them, according to the edge weights. The *ball* $B_v(r)$ of radius r around a vertex v contains all vertices with distance at most r from v.

Definition 1 For a scale $r \in \mathbb{R}_{>0}$, let $\mathcal{P}_{(r,2r]}$ denote the set of all vertex sets of shortest paths with length in $(r, 2r]$. A *shortest path cover* for scale r is a hitting set for $\mathcal{P}_{(r,2r]}$, i.e., a set $\mathrm{SPC}(r) \subseteq V$ such that $|\mathrm{SPC}(r) \cap P| \neq \emptyset$ for all $P \in \mathcal{P}_{(r,2r]}$. The vertices of $\mathrm{SPC}(r)$ are the *hubs* for scale r. A shortest path cover $\mathrm{SPC}(r)$ is *locally h-sparse*, if $|\mathrm{SPC}(r) \cap B_v(2r)| \leq h$ for all vertices $v \in V$. The *highway dimension* of G is the smallest integer h such that there is a locally h-sparse shortest path cover $\mathrm{SPC}(r)$ for every scale $r \in \mathbb{R}_{>0}$ in G.

[1] A metric is said to have *doubling dimension d* if for all $r > 0$ every ball of radius r can be covered by at most 2^d balls of half the radius $r/2$.

[2] It is often assumed that all shortest paths are unique when defining the highway dimension, since this allows good polynomial approximations of this graph parameter [1]. In this work however, we do not rely on these approximations, and thus do not require uniqueness of shortest paths.

The algorithmic consequences of this graph parameter were originally studied in the context of road networks [1–3], which have been shown to have small highway dimension [9]. Road networks are generally non-planar due to overpasses and tunnels, and are also not Euclidean due to different driving or transmission speeds. This is even more pronounced in public transportation networks, where large stations have many incoming connections and plenty of crossing links, making Euclidean (or more generally low doubling) and planar metrics unsuitable as models. Here the highway dimension is better suited, since longer connections are serviced by larger and sparser stations (such as train stations and airports) that can act as hubs. For example, in so-called hub-and-spoke networks of air traffic networks, smaller airports provide connections to close-by larger airports, from which long distance connections are available. This implies a star-like structure of such transportation networks (cf. Sect. 1.3).

The main question posed in this paper is whether the structure of graphs with low highway dimension admits PTASs for problems such as TSP and STP, similar to Euclidean or planar instances. It was shown that *quasi-polynomial time approximation schemes (QPTASs)* exist for these problems [26], i.e., $(1+\varepsilon)$-approximation algorithms with runtime $2^{\mathrm{polylog}(n)}$ assuming that ε and the highway dimension of the input graph are constants. However it was left open whether this can be improved to polynomial time. This paper answers this open question for graphs of highway dimension 1, by giving both lower and upper bounds on the algorithmic complexity for TSP and STP.

1.1 Our Results

Our main result concerns graphs of the smallest possible highway dimension, and shows that for these *fully polynomial time approximation schemes (FPTASs)* exist, i.e., a $(1+\varepsilon)$-approximation can be computed in time polynomial in both the input size and $1/\varepsilon$. Thus at least for this restricted case we obtain a significant improvement over the previously known QPTAS [26]. In particular, the following theorem implies that TSP and STP are polynomial-time solvable on graphs with polynomially bounded edge weights, if the highway dimension is 1.

Theorem 1 *Both Travelling Salesperson and Steiner Tree admit an FPTAS on graphs with highway dimension* 1.

We also show that both the TSP and the STP problem are non-trivial on graphs highway dimension 1, since they are still NP-hard even on this restricted case. This answers an open problem in [26] about the hardness of TSP and STP on graphs of constant highway dimension. Interestingly, together with Theorem 1 this implies that both TSP and STP are *weakly* NP-hard on graphs of highway dimension 1, since strongly NP-hard problems do not admit FPTASs [56], unless P = NP. This is in contrast to planar graphs or Euclidean metrics, for which TSP and STP are strongly NP-hard [29, 30, 50].

Theorem 2 *The* TRAVELLING SALESPERSON *problem is weakly* NP*-hard on graphs with highway dimension* 1.

Theorem 3 *The* STEINER TREE *problem is weakly* NP*-hard on graphs with highway dimension* 1.

1.2 Techniques

We present a step towards a better understanding of low highway dimension graphs by giving new structural insights on graphs of highway dimension 1. It is not hard to find examples of (weighted) complete graphs on arbitrary many vertices with highway dimension 1 (cf. [26]), and thus the class of these graphs does not exclude any minors. Nevertheless, it was suggested in [26] that the *treewidth* of low highway dimension graphs might be bounded polylogarithmically in terms of the *aspect ratio* α, which is the maximum distance divided by the minimum distance between any two vertices of the input graph.

Definition 2 A *tree decomposition* of a graph $G = (V, E)$ is a tree D where each node v is labelled with a bag $X_v \subseteq V$ of vertices of G, such that the following holds:

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

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

As suggested in [26], one may hope to prove that the treewidth of any graph of highway dimension h is, say, $O(h \log(\alpha))$. As argued in Sect. 5, under standard complexity assumptions unfortunately such a bound is generally impossible. In contrast to this, our main structural insight on graphs of highway dimension 1 is that they have treewidth $O(\log \alpha)$. This implies FPTASs for TSP and STP, since we may reduce the aspect ratio of any graph with n vertices to $O(n/\varepsilon)$ and then use algorithms by Bodlaender et al. [15] to compute optimum solutions to TSP and STP in graphs of treewidth t in $2^{O(t)}n$ time. Since reducing the aspect ratio distorts the solution by a factor of $1 + \varepsilon$, this results in an approximation scheme. Although these are fairly standard techniques for metrics (cf. [26]), in our case we need to take special care, since we need to bound the treewidth of the graphs resulting from this reduction, which the standard techniques do not guarantee.

To bound the treewidth of graphs of highway dimension 1, on a high level our technique resembles one that has been used for embeddings of low highway dimension graphs [26] but also low doubling metrics [54] into bounded treewidth graphs. That is, we first find a decomposition of a given graph into a laminar family of clusters (for Euclidean and low doubling metrics this is referred to as a *split-tree* *decomposition*). The laminarity naturally gives rise to a tree structure, and for each cluster we then identify a bag for the corresponding node in the tree. In previous work, the bags were found by selecting a small subset of the vertices contained in the cluster (which approximate the distances well enough to obtain an embedding with small distortion). The novelty of our approach is to shift the focus from the interior of the clusters to the exterior. That is, a bag will contain so-called *interface points* of the cluster, which are vertices that connect the cluster vertices to the outside. In particular, these interface points may not be part of the cluster itself but instead can lie in completely different clusters.

Very recently, this novel shift in perspective has been successfully utilized to obtain the first PTASs for clustering problems such as k -Median and k -Means in graphs of low highway dimension [28]. Whether this will also lead to PTASs for TSP and STP in graphs of bounded highway dimension larger than 1 remains an intriguing open problem. More generally, it remains an open problem to understand the structure of graphs of constant highway dimension larger than 1.

Throughout this paper we use standard graph theoretic notions. For definitions see [23].

1.3 Related Work

The Travelling Salesperson problem (TSP) is among Karp's initial list of 21 NP -complete problems [38]. For general metric instances, for several decades the best known approximation algorithm was due to Christophides [22], which computes a solution with cost at most 3/2 times the LP value. Very recently, an algorithm computing a $(3/2 - \varepsilon)$-approximation (for some $\varepsilon > 10^{-36}$) has been found by Karlin et al. [37]. For unweighted instances, the best known approximation guarantee is 7/5 and is due to Sebő and Vygen [53]. In general the problem is APX-hard [39, 44, 51]. For geometric instances where the nodes are points in $\mathbb{R}^d$ and distances are given by some l_p-norm, there exists a PTAS [4, 49] for fixed d. When $d = \log n$, the problem is APX-hard [55]. Krauthgamer and Lee [43] generalized the PTAS to hyperbolic space. Grigni et al. [32] gave a PTAS for unweighted planar graphs which was later generalized by Arora et al. [5] to the weighted case. For improvements of the running time, see Klein [41].

The Steiner Tree problem (STP) is contained in Karp's list of NP-complete problems as well [38]. The best approximation algorithm known for general metric instances is due to Byrka et al. [18] and computes a solution with cost at most $\ln(4) + \epsilon < 1.39$ times that of an LP relaxation. Their algorithm improved upon previous results by, e.g., Robins and Zelikovsky [52] and Hougardy and Prömel [36]. Also the STP is APX-hard [20] in general. For Euclidean distances and nodes in $\mathbb{R}^d$ with d constant there is a PTAS due to Arora [4]. For $d = \log|R| / \log\log|R|$ where R is the terminal set, the problem is APX-hard [55]. For planar graphs, there is a PTAS for STP [17], and even for the more general Steiner Forest problem for graphs with bounded genus [10]. Note that STP remains NP-complete for planar graphs [30].

It is worth mentioning that alternate definitions of the highway dimension exist.[3] In particular, in a follow-up paper to [3], Abraham et al. [2] define a version of the highway dimension, which implies that the graphs also have bounded doubling dimension. A related model for transportation networks was given by Kosowski and Viennot [42] via the so-called *skeleton dimension*, which also implies bounded doubling dimension. Hence for these definitions, Bartal et al. [7] already provide a PTAS for TSP. The highway dimension definition used here (cf. Definition 1) on the other hand allows for metrics of large doubling dimension as noted by Abraham et al. [3]: a star has highway dimension 1 (by using the center vertex to hit all paths), but its doubling dimension is unbounded. While it may be reasonable to assume that road networks (which are the main concern in the works of Abraham et al. [1–3]) have low doubling dimension, there are metrics modelling transportation networks for which it can be argued that the doubling dimension is large, while the highway dimension should be small. These settings are better captured by Definition 1. For instance, the so-called hub-and-spoke networks that can typically be seen in air traffic networks are star-like networks and are unlikely to have small doubling dimension while still having very small highway dimension close to 1. Thus in these examples it is reasonable to assume that the doubling dimension is a lot larger than the highway dimension.

Feldmann et al. [26] showed that graphs with low highway dimension can be embedded into graphs with low treewidth. This embedding gives rise to a QPTAS for both TSP and STP but also other problems. However, the result in [26] is only valid for a less general definition of the highway dimension from [1], i.e., there are graphs which have constant highway dimension according to Definition 1 but for which the algorithm of [26] cannot be applied. For the less general definition from [1], Becker et al. [11] give a PTAS for Bounded-Capacity Vehicle Routing in graphs of bounded highway dimension. Also the k -Center problem has been studied on graphs of bounded highway dimension, both for the less general definition [11] and the more general one used here [25, 27].

2 Structure of Graphs with Highway Dimension 1

In this section, we analyse the structure of graphs with highway dimension 1. To this end, let us fix a graph G with highway dimension 1 and a shortest path cover $\mathrm{SPC}(r)$ for each scale $r \in \mathbb{R}^+$. As a preprocessing, we remove edges that are longer than the shortest path between their endpoints. This implies that the triangle inequality holds, i.e., the graph has no short cuts.

We begin by analysing the structure of the graph $G_{\leq 2r}$, which is spanned by all edges of the input graph G of length at most $2r$, i.e., we remove all edges of length more than $2r$ from G to obtain $G_{\leq 2r}$. If G has highway dimension 1 it exhibits the following key property.

[3] See [26, Section 9] and [14] for detailed discussions on different definitions of the highway dimension.

Lemma 1 *Let G be a metric graph with highway dimension* 1, $r \in \mathbb{R}^+$ *a scale, and* SPC (r) *a shortest path cover for scale r. Then, every connected component of $G_{\leq 2r}$ contains at most one hub of* SPC (r).

Proof For the sake of contradiction, let $r \in \mathbb{R}^+$ and let $x, y \in$ SPC (r) be a closest pair of distinct hubs in some connected component of $G_{\leq 2r}$. Furthermore, let P be a shortest path in $G_{\leq 2r}$ between x and y, which hence uses only edges of length at most $2r$. (Note that P need not be a shortest path between x and y in G.) In particular, there is no other hub from SPC $(r) \setminus \{x, y\}$ along P. This implies that every edge of P that is not incident to either x or y must be of length at most r, since otherwise the edge would be a shortest path of length $(r, 2r]$ between its endpoints (using that G is metric) contradicting the fact that SPC (r) is a shortest path cover for scale r.

Since the highway dimension of G is 1, any ball $B_w(2r)$ around a vertex $w \in V(P)$ contains at most one of the hubs $x, y \in$ SPC (r). Let $x', y' \in P$ be the vertices adjacent to x and y along P, respectively. Since the length of the edge $\{x, x'\}$ is at most $2r$, the ball $B_{x'}(2r)$ must contain x and, by the observation above, it cannot contain y (in particular $\{x, y\}$ is not an edge). Symmetrically, the ball $B_{y'}(2r)$ contains y but not x. Consequently, $x' \neq y'$ and neither of these two vertices can be a hub of scale r, i.e., the path P contains at least two vertices different from x and y.

Let $V_x = \{w \in V : \operatorname{dist}(x, w) < \operatorname{dist}(y, w)\}$ contain all vertices closer to x than to y, where $\operatorname{dist}(\cdot, \cdot)$ refers to the distance in the original graph G. As all edge weights are strictly positive, we have that $\operatorname{dist}(x, y) > 0$ and thus $y \notin V_x$. Since P starts with vertex $x \in V_x$ and ends with vertex $y \notin V_x$ we deduce that there is an edge $\{u, v\}$ of P such that $u \in V_x$ and $v \notin V_x$. In particular, $\operatorname{dist}(x, u) < \operatorname{dist}(y, u)$ and $\operatorname{dist}(y, v) \leq \operatorname{dist}(x, v)$. We must have $\{u, v\} \neq \{y', y\}$, since otherwise $\operatorname{dist}(x, y') < \operatorname{dist}(y, y') \leq 2r$ and hence $B_{y'}(2r)$ would contain x. Similarly, we have $\{u, v\} \neq \{x, x'\}$, since otherwise $B_{x'}(2r)$ would contain y. Note that, by definition, $u \neq y$ and $v \neq x$, and hence $x, y \notin \{u, v\}$. Consequently, since every edge of P not incident to either x or y must have length at most r, we conclude that $\{u, v\}$ has length at most r.

Finally, consider the scale $r' \in \mathbb{R}^+$, defined such that $2r' = \operatorname{dist}(x, u) + \operatorname{dist}(u, v)$. W.l.o.g., assume that $\operatorname{dist}(x, u) \leq \operatorname{dist}(v, y)$ (otherwise consider scale $2r' = \operatorname{dist}(y, v) + \operatorname{dist}(u, v)$ and the ball $B_u(2r')$ in the following argument). Let Q and Q' denote shortest paths between x, u and v, y in G, respectively. Then the ball $B_v(2r')$ around v contains Q by definition of r'. From $\operatorname{dist}(y, v) \leq \operatorname{dist}(x, v) \leq \operatorname{dist}(x, u) + \operatorname{dist}(u, v) = 2r'$ it follows that $B_v(2r')$ contains Q' as well. Also, $\operatorname{dist}(y, v) \leq \operatorname{dist}(x, v)$ means that $B_v(2r)$ cannot contain x, and hence $2r' = \operatorname{dist}(x, u) + \operatorname{dist}(u, v) \geq \operatorname{dist}(x, v) > 2r$, which implies $r' > r$. Our earlier observation that $\operatorname{dist}(u, v) \leq r$ with $r < r'$ yields $\operatorname{dist}(v, y) \geq \operatorname{dist}(x, u) = 2r' - \operatorname{dist}(u, v) > r'$. In other words, the lengths of both paths Q and Q' are in $(r', 2r']$, and so they both need to contain a hub of SPC (r'). However, by definition of u, v, the paths Q and Q' are vertex disjoint, which means that the ball $B_v(2r')$, which contains Q and Q', also contains at least two hubs from SPC (r'). This is a contradiction with G having highway dimension 1. □

Given a graph G, we now consider graphs $G_{\leq 2r}$ for exponentially growing scales. In particular, for any integer $i \geq 0$ we define the scale $r_i = 2^i$ and call a connected component of $G_{\leq 2r_i}$ a *level-i component*. Note that the level-i components partition the graph G, and that the level-i components are a *refinement* of the level-$(i+1)$ components, i.e., every level-i component is contained in some level-$(i+1)$ component. W.l.o.g., we scale the edge weights of the graph such that $\min_{e \in E} w(e) = 3$, so that there are no edges on level 0, and every level-0 component is a singleton. Recall that $\alpha = \frac{\max_{u \neq v} \operatorname{dist}(u,v)}{\min_{u \neq v} \operatorname{dist}(u,v)} = \frac{\max_{u \neq v} \operatorname{dist}(u,v)}{3}$ is the aspect ratio of G. In our applications we may assume that G is connected, so that there is exactly one level-$(1 + \lceil \log_2(\alpha) \rceil)$ component containing all of G.

Since every edge is a shortest path between its endpoints, every edge $e = \{u, v\}$ that connects a vertex u of a level-i component C with a vertex v outside C is hit by a hub of $\textsc{spc}\,(r_j)$, where j is the level for which $w(e) \in (r_j, 2r_j]$. Moreover, since v lies outside C, we have $w(e) > 2r_i$ and, thus, $j \geq i + 1$. The following definition captures the set of the hubs through which edges can possibly leave C.

Definition 3 Let C be a level-i component of G. We define the set of *interface points* of C as $I_C := \bigcup_{j \geq i} \{u \in \textsc{spc}\,(r_j) : \operatorname{dist}_C(u) \leq 2r_j\}$, where $\operatorname{dist}_C(u)$ denotes the minimum distance from u to a vertex in C (if $u \in C$, $\operatorname{dist}_C(u) = 0$).

Note that, for technical reasons, we explicitly add every hub at level i of a component to its set of interface points as well, even if such a hub does not connect the component with any vertex outside at distance more than $2r_i$.

Lemma 2 *If G has highway dimension* 1, *then each interface I_C of a level-i component C contains at most one hub for each level $j \geq i$.*

Proof Assume that there are two hubs $u, v \in \textsc{spc}\,(r_j)$ in I_C, and recall that we preprocessed the graph so that the triangle inequality holds. Then u and v must be contained in the same level-j component C', since u and v are connected to C with edges of length at most $2r_j$ (or are contained in C) and $C \subseteq C'$. This contradicts Lemma 1. □

Using level-i components and their interface points we can prove that the treewidth of a graph with highway dimension 1 is bounded in terms of the aspect ratio.

Lemma 3 *If a graph G has highway dimension* 1 *and aspect ratio α, its treewidth is at most* $1 + \lceil \log_2(\alpha) \rceil$.

Proof The tree decomposition of G is given by the refinement property of level-i components. That is, let D be a tree that contains a node v_C for every level-i component C for all levels $0 \leq i \leq 1 + \lceil \log_2(\alpha) \rceil$. For every node v_C we add an edge in D to node $v_{C'}$, if C is a level-i component, C' is a level-$(i+1)$ component, and $C \subseteq C'$. The bag X_C for node v_C contains the interface points I_C. For a level-0 component C the bag X_C additionally contains the single vertex u contained in C.

Clearly, the tree decomposition has Property (a) of Definition 2, since the level-0 components partition the vertices of G and every vertex of G is contained in a bag X_C corresponding to a level-0 component C. Also, Property (b) is given by the bags X_C for level-0 components C, since for every edge e of G one of its endpoints u is a hub of $\textsc{spc}(r_i)$ where i is such that $w(e) \in (r_i, 2r_i]$, and the other endpoint w is contained in a level-0 component C, for which X_C contains u and w.

For Property (c), first consider a vertex u of G, which is not contained in any set of interface points for any level-i component and any $0 \leq i \leq \log_2(\alpha)$. Such a vertex only appears in the bag X_C for the level-0 component C containing u, and thus the node v_C for which the bag contains u trivially induces a connected subtree of D.

Any other vertex u of G is an interface point. Let i be the highest level for which $u \in I_C$ for some level-i component C. We claim that $u \in C$, which implies that C is the unique level-i component containing u in its interface. To show our claim, assume $u \notin C$. Then, by definition, I_C contains u because $u \in \textsc{spc}(r_j)$ for some $j \geq i$ and u has some neighbour at distance at most $2r_j$ in C. Since we preprocessed the graph such that every edge is a shortest path between its endpoints, this means that there must be an edge $e = \{u, v\}$ with $w(e) \in (r_j, 2r_j]$ and $v \in C$. Since $u \notin C$, we have $i < j$. Let C' be the unique level-j component with $C \subseteq C'$. Then, by definition, $u \in I_{C'}$, which contradicts the maximality of i. This proves our claim and shows that the highest level component C with $u \in X_C$ is uniquely defined. Moreover, we obtain $u \in \textsc{spc}(r_i)$.

Now consider a level-i' component C' with $i' < i$, such that $u \in X_{C'}$, and let C'' be the unique level-$(i'+1)$ component containing C'. We claim that $u \in X_{C''}$. If $u \in C' \subseteq C''$, then $u \in X_{C''}$, since $u \in \textsc{spc}(r_i)$, $\operatorname{dist}_{C''}(u) = 0 \leq 2r_i$ and $i' + 1 \leq i$. If $u \notin C'$, then $u \in X_{C'}$ implies $u \in I_{C'}$, which means that there must be a vertex $w \in C'$ with $\operatorname{dist}(u, w) \leq 2r_i$. But then $w \in C''$ and thus $\operatorname{dist}_{C''}(u) \leq 2r_i$. Together with $u \in \textsc{spc}(r_i)$, this implies $u \in X_{C''}$, as claimed. Since $v_{C'}$ is a child of $v_{C''}$ in the tree D, it follows inductively that the nodes of D with bags containing u induce a subtree of D with root v_C, which establishes Property (c).

By Lemma 2 each set of interface points contains at most one hub of each level. Since all edges have length at least 3, there are no hubs in $\textsc{spc}(r_0)$ on level 0. This means that each bag of the tree decomposition contains at most $1 + \lceil \log_2(\alpha) \rceil$ interface points. The bags for level-0 components contain one additional vertex. Thus the treewidth of G is at most $1 + \lceil \log_2(\alpha) \rceil$, as claimed. □

An additional property that we will exploit for our algorithms is the following. A (μ, δ)-*net* $N \subseteq V$ is a subset of vertices such that (a) the distance between any two distinct *net points* $u, w \in N$ is more than μ, and (b) for every vertex $v \in V$ there is some net point $w \in N$ at distance at most δ. For graphs of highway dimension 1, we can obtain nets with additional favourable properties, as the next lemma shows.

Lemma 4 *For any graph G of highway dimension 1 and any $r > 0$, there is an $(r, 3r)$-net such that every connected component of $G_{\leq r}$ contains exactly one net point. Moreover this net can be computed in polynomial time.*

Proof We first derive an upper bound of $3r$ for the diameter of any connected component of $G_{\leq r}$. Lemma 1 implies that a connected component C contains at most one hub x of SPC$(r/2)$. By definition, any shortest path in C of length in $(r/2, r]$ must pass through x. We also know that every edge of C has length at most r. Consequently, every edge in C not incident to x must have length at most $r/2$, since each edge constitutes a shortest path between its endpoints. This implies that any shortest path in C that is not hit by x must have length at most $r/2$: if C contains a shortest path P with length more than $r/2$ not containing x we could repeatedly remove edges of length at most $r/2$ from P until we obtain a shortest path of length in $(r/2, r]$ not hit by x, a contradiction. Now consider a shortest path P in G of length more than $r/2$ from some vertex $v \in C$ to x (note that this path may not be entirely contained in C). Let $\{u, w\}$ be the unique edge of P such that $\text{dist}(v, u) \leq r/2$ and $\text{dist}(v, w) > r/2$. If the length of the edge $\{u, w\}$ is at most $r/2$ then $\text{dist}(v, w) \leq r$, and thus $w = x$, since the part of the path from v to w is a shortest path of length in $(r/2, r]$ and thus needs to pass through x. Otherwise the length of the edge $\{u, w\}$ is in the interval $(r/2, r]$, which again implies $w = x$, since the edge must contain x. In either case, $\text{dist}(v, x) \leq 3r/2$. This implies that every vertex in C is at distance at most $3r/2$ from x, and thus the diameter of C is at most $3r$.

To compute the $(r, 3r)$-net, we greedily pick an arbitrary vertex of each connected component of $G_{\leq r}$. As the distances between components of $G_{\leq r}$ is greater than r, and every vertex lies in some component containing a net point, we get the desired distance bounds. Clearly this net can be computed in polynomial time. □

3 Approximation Schemes

In general the aspect ratio of a graph may be exponential in the input size. We need to reduce the aspect ratio α of the input graph $G = (V, E)$ to a polynomial. For both STP and TSP, standard techniques can be used to reduce the aspect ratio to $O(n/\varepsilon)$ when aiming for a $(1 + \varepsilon)$-approximation. This was for instance also used in [26] for low highway dimension graphs, but here we need to take special care not to destroy the structural properties given by Lemma 3 in this process. In particular, we need to reduce the aspect ratio and maintain the fact that the treewidth is bounded.

Therefore, we reduce the aspect ratio of our graphs by the following preprocessing. Both metric TSP and STP admit constant factor approximations in polynomial time using well-known algorithms [18, 22]. We first compute a solution of cost c using a β-approximation algorithm for the problem at hand (TSP or STP). For TSP, the diameter of the graph G clearly is at most $c/2$. For STP we remove every vertex of V that is at distance more than c from any terminal, since such a vertex cannot be part of the optimum solution. After having removed all such vertices in this way, we obtain a graph G of diameter at most $3c$. Thus, in the following, we may assume that our graph G has diameter at most $3c$. We then set $r = \frac{\varepsilon c}{3n}$ in Lemma 4 to obtain a $(\frac{\varepsilon c}{3n}, \frac{\varepsilon c}{n})$-net $N \subseteq V$. As a consequence the metric induced by N (with distances of G) has aspect ratio at most $\frac{3c}{\varepsilon c/(3n)} = O(n/\varepsilon)$, since the minimum distance between any two net points of N is at least $\frac{\varepsilon c}{3n}$ and the maximum distance is at most $3c$. We will exploit this property in the following.

By Lemma 4, each connected component of $G_{\leq \frac{\varepsilon c}{3n}}$ contains exactly one net point of N. Let $\eta : V \mapsto N$ map each vertex of G to the unique net point in the same connected component of $G_{\leq \frac{\varepsilon c}{3n}}$. We define a new graph G' with vertex set $N \subseteq V$ and edge set $\{\{\eta(u), \eta(v)\} : \{u, v\} \in E \wedge \eta(u) \neq \eta(v)\}$. The length of each edge $\{w, w'\}$ of G' is the shortest path distance between w and w' in G. This new graph G' may not have bounded highway dimension, but we claim that it has treewidth $O(\log(n/\varepsilon))$.

Lemma 5 *If G has highway dimension* 1, *the graph G' with vertex set N has treewidth $O(\log(n/\varepsilon))$. Moreover, a tree decomposition for G' of width $O(\log(n/\varepsilon))$ can be computed in polynomial time.*

Proof We construct a tree decomposition D' of G' as follows. In light of Lemma 3, we can compute a tree decomposition D of width at most $1 + \lceil \log_2(\alpha) \rceil$, where α is the aspect ratio of G: for this we need to compute a locally 1-sparse shortest path cover $\text{SPC}(r_i)$ for each level i, which can be done in polynomial time via an XP algorithm [26] if the highway dimension is 1. We then find the level-i components and their interface points, from which the tree decomposition D and its bags can be constructed. Since there are $O(\log \alpha)$ levels and α is at most exponential in the input size (which includes the encoding length of the edge weights), we can compute D in polynomial time.

We construct D' from D by replacing every bag X of D by a new bag $X' = \{\eta(v) : v \in X\}$ containing the net points for the vertices in X. It is not hard to see that Properties (a) and (b) of Definition 2 are fulfilled by D', since they are true for D. For Property (c), note that for any edge $\{u, v\}$ of G, the set of all bags of D that contain u or v form a connected subtree of D. This is because the bags containing u form a connected subtree (Property (c)), the same is true for v, and both these subtrees share at least one node labelled by a bag containing the edge $\{u, v\}$ (Property (b)). Consequently, the set of all bags containing vertices of any connected subgraph of G form a connected subtree. In particular, for any connected component A of $G_{\leq \frac{\varepsilon c}{3n}}$, the set of bags of D containing at least one vertex of A form a connected subtree. This implies Property (c) for D'. Thus, D' is indeed a tree decomposition of G' according to Definition 2. Note that D' can be computed in polynomial time.

To bound the width of D', recall that a bag X of the tree decomposition D of G contains the interface points I_C of a level-i component C, in addition to one more vertex of C on the lowest level $i = 0$. Each interface point is a hub from $\text{SPC}(r_j)$ at some level $j \geq i$ and is at distance at most $2r_j$ from C. In particular, if $2r_i \leq \frac{\varepsilon c}{3n}$ then C is a component of $G_{\leq 2r_i} \subseteq G_{\leq \frac{\varepsilon c}{3n}}$, and all hubs of $I_C \cap \text{SPC}(r_j)$ for which $2r_j \leq \frac{\varepsilon c}{3n}$ lie in the same connected component A of $G_{\leq \frac{\varepsilon c}{3n}}$ as C. These hubs are therefore all mapped to the same net point w in A by η. In addition to w, the bag $X' = \{\eta(v) : v \in X\}$ resulting from X and η contains at most one vertex for every level j such that $2r_j > \frac{\varepsilon c}{3n}$. As $r_j = 2^j$, this condition is equivalent to $j > \log_2(\frac{\varepsilon c}{3n}) - 1$.

As there are $1 + \lceil \log_2(\alpha) \rceil$ levels in total, there are $O(\log(\frac{\alpha n}{\varepsilon c}))$ hubs in X'. This bound is obviously also valid in case $2r_i > \frac{\varepsilon c}{3n}$. We preprocessed the graph G so that its diameter is at most $3c$ and its minimum distance is 3, which implies an aspect ratio α of at most c for G. This means that every bag X' contains $O(\log(n/\varepsilon))$ vertices, and thus the claimed treewidth bound for G' follows. □

We are now ready to prove our main result.

Proof of of Theorem 1 To solve TSP or STP on G we first use the above reduction to obtain G' and its tree decomposition D', and then compute an optimum solution for G'. For TSP, G' is already a valid input instance, but for STP we need to define a terminal set, which simply is $R' = \{\eta(v) \mid v \in R\}$ if R is the terminal set of G. Bodlaender et al. [15] proved that for both TSP and STP there are deterministic algorithms to solve these problems exactly in time $2^{O(t)}n$, given a tree decomposition of the input graph of width t. By Lemma 5 we can thus compute the optimum to G' in time $2^{O(\log(n/\varepsilon))} \cdot n = (n/\varepsilon)^{O(1)}$. Afterwards, we convert the solution for G' back to a solution for G, as follows.

For TSP we may greedily add vertices of V to the tour on N by connecting every vertex $v \in V$ to the net point $\eta(v)$. As the vertices N of G' form a $(\frac{\varepsilon c}{3n}, \frac{\varepsilon c}{n})$-net of V, this incurs an additional cost of at most $2\frac{\varepsilon c}{n}$ per vertex, which sums up to at most $2\varepsilon c$. Let $\textsc{Opt}$ and $\textsc{Opt}'$ denote the costs of the optimum tours in G and G', respectively. We know that $c \leq \beta \cdot \textsc{Opt}$, since we used a β-approximation algorithm to compute c. Furthermore, the optimum tour in G can be converted to a tour in G' of cost at most $\textsc{Opt}$ by short-cutting, due to the triangle inequality. Thus $\textsc{Opt}' \leq \textsc{Opt}$, which means that the cost of the computed tour in G is at most $\textsc{Opt}' + 2\varepsilon c \leq (1 + 2\beta\varepsilon)\textsc{Opt}$.

Similarly, for STP we may greedily connect a terminal v of G to the terminal $\eta(v)$ of G' in the computed Steiner tree in G'. This adds an additional cost of at most $\frac{\varepsilon c}{n}$, which sums up to at most εc. Let now $\textsc{Opt}$ and $\textsc{Opt}'$ be the costs of the optimum Steiner trees in G and G', respectively. We may convert a Steiner tree T in G into a tree T' in G' by using edge $\{\eta(u), \eta(v)\}$ for each edge $\{u, v\}$ of T. Note that the resulting tree T' contains all terminals of G', since $R' = \{\eta(v) \mid v \in R\}$. As the vertices N of G' form a $(\frac{\varepsilon c}{3n}, \frac{\varepsilon c}{n})$-net of V, the cost of T' is at most $\textsc{Opt} + 2\varepsilon c$ if the cost of T is $\textsc{Opt}$ (by the same argument as used for the proof of Lemma 5). As before, we know that $c \leq \beta \cdot \textsc{Opt}$, and thus the cost of the computed Steiner tree in G is at most $\textsc{Opt}' + \varepsilon c \leq \textsc{Opt} + 3\varepsilon c \leq (1 + 3\beta\varepsilon)\textsc{Opt}$.

Hence we obtain FPTASs for both TSP and STP computing $(1 + \varepsilon)$-approximations within a runtime that is polynomial in the input size and $1/\varepsilon$. □

4 Hardness Proofs

In this section, we prove next that both TSP and STP are NP-hard on graphs of highway dimension 1 implying that these problems are weakly NP-hard for these inputs (cf. [56]).

4.1 Hardness of Travelling Salesperson for Highway Dimension 1

We present a reduction from Hamiltonian Cycle which was shown to be NP-complete by Karp [38]. For this problem, we are given an undirected graph $G = (V, E)$ and the task is to decide whether G admits a Hamiltonian cycle, i.e., a cycle that visits each vertex exactly once.

Proof of Theorem 2 we reduce from Hamiltonian Cycle, let an undirected graph G be given. For the reduction, we endow the edges of G with a weight function $w : E \to \mathbb{R}$ such that the resulting weighted graph has highway dimension 1 and the value of an optimal TSP tour decides the Hamiltonicity of G. Specifically, let $V = \{v_0, \dots, v_{n-1}\}$, and for each edge $e = \{v_i, v_j\} \in E$, we let $w(e) = 11^i + 11^j$. □

Claim 1 *The resulting edge-weighted graph G has highway dimension* 1.

Proof Fix an arbitrary scale $r > 0$. Since there is no edge with weight less than 12, we may assume without loss of generality that $r \geq 6$. Let $i = \lceil \log_{11}(r/5) \rceil \geq 0$ and define $\text{SPC}(r) = \{v_i\}$. It is left to show that every shortest path with length in $(r, 2r]$ uses v_i. Note that any edge incident to a vertex v_j with $j \geq i + 1$ has length at least $11^{i+1} + 1 \geq 11r/5 + 1 > 2r$. On the other hand, consider dividing each edge $e = \{v_i, v_j\}$ into two half-edges with length 11^i and 11^j, respectively. Then, visiting vertex v_i involves using two half-edges with length 11^i each, one for entering v_i and one for leaving v_i. Thus all shortest paths that do not use any v_j with $j \geq i$ have length at most

$$\sum_{j=0}^{i-1} 2 \cdot 11^j,$$

since each shortest path is simple and thus visits the vertices $v_0, v_1, \dots, v_{i-1}$ at most once, and every visit of a vertex v_j with $j \in \{0, \dots, i-1\}$ contributes at most $2 \cdot 11^i$ to the cost of the path. We obtain

$$2\sum_{j=0}^{i-1} 11^j = 2\frac{11^i - 1}{11 - 1} = \frac{11^i}{5} - \frac{1}{5} < \frac{11^i}{5} \leq \frac{11^{\log_{11}(r/5)+1}}{5} = \frac{11}{25}r < r.$$

Thus, the only shortest paths with length in $(r, 2r]$ are those passing through v_i, as claimed. □

To finish the reduction, we show that G is Hamiltonian if and only if the optimal TSP tour in the weighted version of G has length $2\sum_{i=0}^{n-1} 11^i = \frac{11^n}{5} - \frac{1}{5}$. First note that every tour needs to visit every vertex once. For each $i \in \{0, \dots, n-1\}$, a visit of vertex v_i contributes $2 \cdot 11^i$ to the cost of the tour. Thus, the cost of an optimal TSP tour is at least $\frac{11^n}{5} - \frac{1}{5}$, and this cost is realized if G is Hamiltonian. If, on the other hand, G is not Hamiltonian, an optimal TSP tour needs to visit at least one

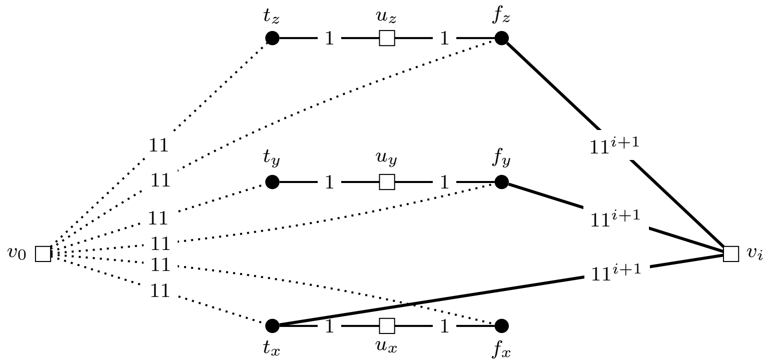

Fig. 1 Illustration of the part of the construction involving vertices x, y, z and a clause $C_i = (x \vee \bar{y} \vee \bar{z})$. Terminals are marked as boxes

vertex at least twice, so that in this case the length of the optimal TSP tour is at least $\frac{11^n}{5} - \frac{1}{5} + 2$. □

4.2 Hardness of Steiner Tree for Highway Dimension 1

We present a reduction from the NP-hard satisfiability problem (SAT) [31], in which a Boolean formula φ in conjunctive normal form is given, and a satisfying assignment of its variables needs to be found.

Proof of of Theorem 3 For a given SAT formula φ with k variables and ℓ clauses we construct a graph G_φ as follows (cf. Fig. 1). For each variable x we introduce a path $P_x = (t_x, u_x, f_x)$ with two edges of length 1 each. The vertex u_x is a terminal. Additionally we introduce a terminal v_0, which we call the *root*, and add the edges $\{v_0, t_x\}$ and $\{v_0, f_x\}$ for every variable x. Every edge incident to v_0 has length 11. For each clause C_i, where $i \in \{1, \ldots, \ell\}$, we introduce a terminal 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 11^{i+1}. Note that the edges incident to the root v_0 also have length 11^{i+1} for $i = 0$.

Claim 2 *The constructed graph* G_φ *has highway dimension* 1.

Proof Fix a scale $r > 0$. If $r \leq 5$ then the shortest path cover SPC (r) only needs to hit shortest paths of length at most $2r \leq 10$. Since all edges incident to terminals v_j with $j \in \{0, \ldots, \ell\}$ have length at least 11, any such path contains only edges of paths P_x. Thus it suffices to include all vertices u_x in SPC (r). A ball $B_w(2r)$ of radius $2r \leq 10$ can also only contain some subset of vertices of a single path P_x, or a single vertex

v_j. In the former case the ball contains at most the vertex $u_x \in \text{SPC}(r)$, and in the latter none of $\text{SPC}(r)$.

If $r > 5$, let $i = \lfloor \log_{11}(r/5) \rfloor \geq 0$ and $\text{SPC}(r) = \{v_i\}$. Since there is only one hub, this shortest path cover is locally 1-sparse. Note that any edge incident to a vertex v_j with $j \geq i + 1$ has length at least $11^{i+2} \geq 11r/5 > 2r$. Also, all paths that do not use any v_j with $j \geq i$ have length at most $2 + \sum_{j=0}^{i-1}(2 \cdot 11^{j+1} + 2)$, since such a path can contain at most two edges incident to a vertex v_j with $j \leq i - 1$ 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 shorter than

$$2 + 2\frac{11^{i+1}}{11 - 1} + 2i \leq 3 \cdot 11^i + 2 \cdot 11^i \leq 5 \cdot 11^i \leq r,$$

where the first inequality holds since $i + 1 \leq 11^i$ whenever $i \geq 0$. Hence the only paths that need to be hit by hubs on scale r are those passing through v_i, which is a hub of $\text{SPC}(r)$. □

To finish the reduction, we claim that there is a satisfying assignment for φ if and only if there is a Steiner tree T for G_φ with cost at most $12k + \sum_{i=1}^{\ell} 11^{i+1}$. If there is a satisfying assignment for φ, then the tree T contains the edges $\{u_x, t_x\}$ and $\{v_0, t_x\}$ for variables x that are set to true, and the edges $\{u_x, f_x\}$ and $\{v_0, f_x\}$ for variables x that are set to false. This connects every terminal u_x with the root v_0, and the cost of these edges is $12k$. For every terminal v_i where $i \geq 1$ we can now add the edge $\{v_i, s_x\}$ for $s_x \in \{t_x, f_x\}$ that corresponds to a literal of C_i that is true in the satisfying assignment. Since this Steiner vertex s_x is connected to the root v_0, we obtain a Steiner tree T. The latter edges add another $\sum_{i=1}^{\ell} 11^{i+1}$ to the solution cost, and thus the total cost is as claimed.

Conversely, consider a minimum cost Steiner tree T in G_φ. Note that for any terminal u_x the tree must contain an incident edge of cost 1, while for any terminal v_i with $i \geq 1$ the tree must contain an incident edge of cost 11^{i+1}. This adds up to a cost of $k + \sum_{i=1}^{\ell} 11^{i+1}$. Assume that there is some variable x such that T contains neither $\{v_0, t_x\}$ nor $\{v_0, f_x\}$. This means that in T the terminal u_x is connected to the root v_0 through an edge $\{v_i, s_x\}$ for $s_x \in \{t_x, f_x\}$ and some $i \geq 1$. The edge $\{v_0, s_x\}$ forms a fundamental cycle with the tree T, which however has a shorter length of 11 compared to the edge $\{v_i, s_x\}$, which has length 11^{i+1}. Thus removing $\{v_0, s_x\}$ and adding $\{v_i, s_x\}$ instead, would yield a cheaper Steiner tree. As this would contradict that T has minimum cost, T contains at least one of the edges $\{v_0, t_x\}$ and $\{v_0, f_x\}$ for every variable x. This adds another $11k$ to the cost, so that T costs at least $12k + \sum_{i=1}^{\ell} 11^{i+1}$.

If we assume that $12k + \sum_{i=1}^{\ell} 11^{i+1}$ is also an upper bound on the cost of T, by the above observations the tree T contains exactly one edge incident to every terminal u_x and v_i for $i \geq 1$, and exactly k edges incident to v_0. Furthermore, for every variable x the latter edges contain exactly one of $\{v_0, t_x\}$ and $\{v_0, f_x\}$. Thus T encodes a satisfying assignment for φ, as follows. For every edge $\{v_0, t_x\}$ we may set x to true, and for every edge $\{v_0, f_x\}$ we may set x to false. For every clause C_i the corresponding terminal v_i connects through one of the Steiner vertices $s_x \in \{t_x, f_x\}$ of a corresponding literal contained in C_i. The only incident vertices to s_x in G_φ are some terminals v_j, the terminal u_x, and the root v_0. As each v_j and also u_x only has one incident edge contained in the tree T, the tree must contain the edge $\{v_0, s_x\}$ so that the root can be reached from s_x in T. Hence s_x corresponds to a literal that is true in C_i. Using Lemma 2, which bounds the highway dimension of G_φ, we obtain Theorem 3. □

5 Conclusions

We showed that, somewhat surprisingly, graphs of highway dimension 1 exhibit a rich combinatorial structure. On one hand, it was already known [26] that these graphs do not exclude any minors and thus their treewidth is unbounded. Here we additionally showed that STP and TSP are weakly NP-hard on such graphs, further confirming that these graphs have non-trivial properties. On the other hand, we proved in Lemma 3 that the treewidth of a graph of highway dimension 1 is logarithmically bounded in the aspect ratio α. This in turn can be exploited to obtain a very efficient FPTAS for both STP and TSP.

At this point one may wonder whether it is possible to generalize Lemma 3 to larger values of the highway dimension. In particular, in [26] it was suggested that the treewidth of a graph of highway dimension h might be bounded by, say, $O(h \log(\alpha))$. However such a bound is highly unlikely in general, since it would have the following consequence for the k -Center problem, for which k vertices (centers) need to be selected in a graph such that the maximum distance of any vertex to its closest center is minimized. It was shown in [25] that it is NP-hard to compute a $(2-\varepsilon)$-approximation for k -Center on graphs of highway dimension $O(\log^2 n)$, for any $\varepsilon > 0$. Given such a graph, the same preprocessing of Sect. 3 could be used to derive an analogue of Lemma 5, i.e., a graph G' of treewidth $O(\text{polylog}(n/\varepsilon))$ could be computed for the net N. Moreover, a 2-approximation for k -Center can be computed in polynomial time on any graph [35], and if the input has treewidth t a $(1+\varepsilon)$ -approximation can be computed in $(t/\varepsilon)^{O(t)} n^{O(1)}$ time [40]. Using the same arguments to prove Theorem 1 for STP and TSP, it would now be possible to compute a $(1+\varepsilon)$-approximation for k -Center in quasi-polynomial time (cf. [27]). That is, we would obtain a QPTAS for graphs of highway dimension $O(\log^2 n)$, which is highly unlikely given that computing a $(2-\varepsilon)$-approximation is NP-hard on such graphs.

The above argument in fact rules out any bound of $(h \log \alpha)^{O(1)}$ for graphs of highway dimension h and aspect ratio α, unless NP-hard problems admit quasi-polynomial time algorithms. In fact, we conjecture that the k -Center problem is NP-hard to approximate within a factor of $2-\varepsilon$ for graphs of constant highway dimension (for some constant larger than 1). If this is true, then the above argument even rules out a treewidth bound of $(\log \alpha)^{f(h)}$ for any function f. Thus, in order to answer the open problem of [26] and obtain a PTAS for graphs of constant highway dimension, a different approach seems to be needed.

Acknowledgements The first author is supported by the 'Excellence Initiative' of the German Federal and State Governments and the Graduate School CE at TU Darmstadt. The second author is 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). The third author is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy

The Berlin Mathematics Research Center MATH+ (EXC-2046/1, Project ID: 390685689). The fourth author is supported by the Discovery Grant Program of the Natural Sciences and Engineering Research Council of Canada. A preliminary version (extended abstract) of this paper appeared at WG 2019 [24]

References

1. Abraham, I., Delling, D., Fiat, A., Goldberg, A.V., Werneck, R.F.: VC-dimension and shortest path algorithms. In: Proceedings of 28th International Colloquium on Automata, Languages, and Programming (ICALP), pp. 690–699 (2011). https://doi.org/10.1007/978-3-642-22006-7_58
2. Abraham, I., Delling, D., Fiat, A., Goldberg, A.V., Werneck, R.F.: Highway dimension and provably efficient shortest path algorithms. J. ACM (2016). https://doi.org/10.1145/2985473
3. Abraham, I., Fiat, A., Goldberg, A.V., Werneck, R.F.: Highway dimension, shortest paths, and provably efficient algorithms. In: Proceedings of 21st Annual ACM-SIAM Symposium Discrete Algorithms (SODA), pp. 782–793 (2010). https://doi.org/10.1137/1.9781611973075.64
4. Arora, S.: Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems. J. ACM **45**(5), 753–782 (1998). https://doi.org/10.1145/290179.290180
5. Arora, S., Grigni, M., Karger, D.R., Klein, P.N., Woloszyn, A.: A polynomial-time approximation scheme for weighted planar graph TSP. In: Proceedings of 9th Annual ACM-SIAM Symposium Discrete Algorithms (SODA), pp. 33–41 (1998)
6. Arora, S., Lund, C., Motwani, R., Sudan, M., Szegedy, M.: Proof verification and hardness of approximation problems. In: Proceedings of 33rd Annual IEEE Symposium on Foundations Computer Science (FOCS), pp. 14–23 (1992). https://doi.org/10.1145/278298.278306
7. Bartal, Y., Gottlieb, L.A., Krauthgamer, R.: The traveling salesman problem: low-dimensionality implies a polynomial time approximation scheme. In: Proceedings of 44th Annual ACM Symposium on Theory Computing (STOC), pp. 663–672 (2012). https://doi.org/10.1145/2213977.2214038
8. Bast, H., Funke, S., Matijevic, D.: Ultrafast shortest-path queries via transit nodes. In: The Shortest Path Problem: Ninth DIMACS Implementation Challenge **74**, 175–192 (2009)
9. Bast, H., Funke, S., Matijevic, D., Sanders, P., Schultes, D.: In transit to constant time shortest-path queries in road networks. In: Proceedings of 9th Workshop Algorithm Engineering and Experiments (ALENEX) (2007). https://doi.org/10.1137/1.9781611972870.5
10. Bateni, M., Hajiaghayi, M.T., Marx, D.: Approximation schemes for Steiner forest on planar graphs and graphs of bounded treewidth. J. ACM **58**, 21:1–21:37 (2011). https://doi.org/10.1145/2027216.2027219
11. Becker, A., Klein, P.N., Saulpic, D.: Polynomial-time approximation schemes for k-center, k-median, and capacitated vehicle routing in bounded highway dimension. In: Proceedings of 26th European Symposium Algorithms (ESA), pp. 8:1–8:15 (2018). https://doi.org/10.4230/LIPIcs.ESA.2018.8
12. Bern, M., Plassmann, P.: The Steiner problem with edge lengths 1 and 2. Inf. Process. Lett. **32**, 171–176 (1989). https://doi.org/10.1016/0020-0190(89)90039-2
13. Bland, R., Shallcross, D.: Large traveling salesman problems arising from experiments in X-ray crystallography: a preliminary report on computation. Oper. Res. Lett. **8**, 125–128 (1989). https://doi.org/10.1016/0167-6377(89)90037-0
14. Blum, J.: Hierarchy of transportation network parameters and hardness results (2019). arXiv:1905.11166
15. Bodlaender, H.L., Cygan, M., Kratsch, S., Nederlof, J.: Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. Inf. Comput. **243**, 86–111 (2015). https://doi.org/10.1016/j.ic.2014.12.008
16. Borndörfer, R., Neumann, M., Pfetsch, M.E.: The line connectivity problem. In: Fleischmann, B., Borgwardt, K.H., Klein, R., Tuma, A.: (eds.) Operations Research Proceedings, pp. 557–562 (2009). https://doi.org/10.1007/978-3-642-00142-0_90
17. Borradaile, G., Kenyon-Mathieu, C., Klein, P.: A polynomial-time approximation scheme for Steiner tree in planar graphs. In: Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1285–1294 (2007). https://doi.org/10.1145/1541885.1541892

18. Byrka, J., Grandoni, F., Rothvoß, T., Sanità, L.: An improved LP-based approximation for Steiner tree. In: Proceedings of the 42nd Annual ACM Symposium on Theory Computing (STOC), pp. 583–592 (2010). https://doi.org/10.1145/1806689.1806769
19. Chen, C.Y., Grauman, K.: Efficient activity detection in untrimmed video with max-subgraph search. IEEE Trans. Pattern Anal. Mach. Intell. **39**, 908–921 (2018). https://doi.org/10.1109/TPAMI.2016.2564404
20. Chlebík, M., Chlebíková, J.: The Steiner tree problem on graphs: Inapproximability results. Theor. Comput. Sci. **406**, 207–214 (2008). https://doi.org/10.1016/j.tcs.2008.06.046
21. Chowdhury, S.A., Shackney, S.E., Heselmeyer-Haddad, K., Ried, T., Schäffer, A.A., Schwartz, R.: Phylogenetic analysis of multiprobe fluorescence in situ hybridization data from tumor cell populations. Bioinformatics **29**, i189–i198 (2013). https://doi.org/10.1093/bioinformatics/btt205
22. Christofides, N.: Worst-case analysis of a new heuristic for the travelling salesman problem. Technical Report 388, Graduate School of Industrial Administration, Carnegie Mellon University (1976). http://www.dtic.mil/dtic/tr/fulltext/u2/a025602.pdf
23. Diestel, R.: Graph Theory. Graduate Texts in Mathematics, vol. 173, 5th edn. Springer, Berlin (2017)
24. Disser, Y., Feldmann, A.E., Klimm, M., Könemann, J.: Travelling on graphs with small highway dimension. In: Proceedings of the 45th International Workshop on Graph-Theoretic Concepts in Computer Science, WG 2019, pp. 175–189 (2019)
25. Feldmann, A.E.: Fixed parameter approximations for k-center problems in low highway dimension graphs. Algorithmica **81**, 1031–1052 (2019). https://doi.org/10.1007/s00453-018-0455-0
26. Feldmann, A.E., Fung, W.S., Könemann, J., Post, I.: A $(1+\varepsilon)$-embedding of low highway dimension graphs into bounded treewidth graphs. SIAM J. Comput. **41**, 1667–1704 (2018). https://doi.org/10.1137/16M1067196
27. Feldmann, A.E., Marx, D.: The parameterized hardness of the k-center problem in transportation networks. In: Proceedings of 16th Scandinavian Symposium and Workshop Algorithm Theory (SWAT), pp. 19:1–19:13 (2018). https://doi.org/10.4230/LIPIcs.SWAT.2018.19
28. Feldmann, A.E., Saulpic, D.: Polynomial time approximation schemes for clustering in low highway dimension graphs. In: Proceedings of the 28th European Symposium on Algorithms (ESA), vol. 173, pp. 46:1–46:22 (2020). arXiv:2006.12897
29. Garey, M.R., Graham, R.L., Johnson, D.S.: The complexity of computing Steiner minimal trees. SIAM J. Appl. Math **32**, 835–859 (1977). https://doi.org/10.1137/0132072
30. Garey, M.R., Johnson, D.S.: The rectilinear Steiner tree problem is NP-complete. SIAM J. Appl. Math. **32**, 826–834 (1977). https://doi.org/10.1137/0132071
31. Garey, M.R., Johnson, D.S.: Computers and intractability. Freeman **vol. 29**, (2002)
32. Grigni, M., E. Koutsoupias, Papadimitriou, C.H.: An approximation scheme for planar graph TSP. In: Proceedings of 36th Annual IEEE Symposium on Foundations Computer Science (FOCS), pp. 640–645 (1995). https://doi.org/10.1109/SFCS.1995.492665
33. Grötschel, M., Holland, O.: Solution of large-scale symmetric travelling salesman problems. Math. Program. **51**, 141–202 (1991). https://doi.org/10.1007/BF01586932
34. Held, S., Korte, B., Rautenbach, D., Vygen, J.: Combinatorial optimization in VLSI design. In: Chvatal, V. (ed.) Combinatorial Optimization: Methods and Applications, pp. 33–96. IOS Press, Amsterdam (2011)
35. Hochbaum, D.S., Shmoys, D.B.: A unified approach to approximation algorithms for bottleneck problems. J. ACM **33**(3), 533–550 (1986). https://doi.org/10.1145/5925.5933
36. Hougardy, S., Prömel, H.J.: A 1.598 approximation algorithm for the Steiner problem in graphs. In: Proceedings of 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 448–453 (1999)
37. Karlin, A.R., Klein, N., Gharan, S.O.: A (slightly) improved approximation algorithm for metric TSP. arXiv:2007.01409 (2020)
38. Karp, R.M.: Reducibility among combinatorial problems. In: Complexity of Computer Computations pp. 85–103 (1972)
39. Karpinski, M., Lampis, M., Schmied, R.: New inapproximability bounds for TSP. J. Comput. Syst. Sci. **81**, 1665–1677 (2015). https://doi.org/10.1016/j.jcss.2015.06.003
40. Katsikarelis, I., Lampis, M., Paschos, V.T.: Structural parameters, tight bounds, and approximation for (k, r)-center. In: Proceedings of 28th International Symposium on Algorithms Computing (ISAAC), pp. 50:1–50:13 (2017). https://doi.org/10.4230/LIPIcs.ISAAC.2017.50

41. Klein, P.: A linear-time approximation scheme for TSP in undirected planar graphs with edge-weights. SIAM J. Comput. **37**(6), 1926–1952 (2008). https://doi.org/10.1137/060649562
42. Kosowski, A., Viennot, L.: Beyond highway dimension: small distance labels using tree skeletons. In: Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 1462–1478 (2017)
43. Krauthgamer, R., Lee, J.R.: Algorithms on negatively curved spaces. In: Proceedings of 47th Annual IEEE Symposium on Foundations Computing Science (FOCS), pp. 119–132 (2006). https://doi.org/10.1109/FOCS.2006.9
44. Lampis, M.: Improved inapproximability for TSP. Theory Comput. **10**, 217–236 (2014). https://doi.org/10.4086/toc.2014.v010a009
45. Laporte, G., Nobert, Y., Desrochers, M.: Optimal routing under capacity and distance restrictions. Oper. Res. **33**, 1050–1073 (1985). https://doi.org/10.1287/opre.33.5.1050
46. Lenstra, J.K., Rinnooy Kan, A.H.G.: Some simple applications of the traveling salesman problem. Oper. Res. Quart. **26**, 717–33 (1975). https://doi.org/10.2307/3008306
47. Ljubić, I., Weiskirchner, R., Pferschy, U., Klau, G.W., Mutzel, P., Fischetti, M.: An algorithmic framework for the exact solution of the prize-collecting Steiner tree problem. Math. Program. **105**, 427–449 (2006). https://doi.org/10.1007/s10107-005-0660-x
48. Loboda, A.A., Artyomov, M.N., Sergushichev, A.A.: Solving generalized maximum-weight connected subgraph problem for network enrichment analysis. In: Proceedings of the 16th Workshop Algorithms in Bioinformatics (WABI), pp. 210–221 (2016). https://doi.org/10.1007/978-3-319-43681-4_17
49. Mitchell, J.S.B.: Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, *k*-MST, and related problems. SIAM J. Comput. **28**(4), 1298–1309 (1999). https://doi.org/10.1137/S0097539796309764
50. Papadimitriou, C.H.: The Euclidean travelling salesman problem is NP-complete. Theor. Comput. Sci. **4**, 237–244 (1977)
51. Papadimitriou, C.H., Vempala, S.: On the approximability of the traveling salesman problem. Combinatorica **26**, 101–120 (2006). https://doi.org/10.1007/s00493-006-0008-z
52. Robins, G., Zelikovsky, A.: Tighter bounds for graph Steiner tree approximation. SIAM J. Discrete Math. **19**, 122–134 (2005). https://doi.org/10.1137/S0895480101393155
53. Sebő, A., Vygen, J.: Shorter tours by nicer ears: $7/5$-approximation for the graph-TSP, $3/2$ for the path version, and $4/3$ for two-edge-connected subgraphs. Combinatorica (2014). https://doi.org/10.1007/s00493-011-2960-3
54. Talwar, K.: Bypassing the embedding: algorithms for low dimensional metrics. In: Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC), pp. 281–290 (2004). https://doi.org/10.1145/1007352.1007399
55. Trevisan, L.: When Hamming meets Euclid: the approximability of geometric TSP and Steiner tree. SIAM J. Comput. **30**, 475–485 (2000). https://doi.org/10.1137/S0097539799352735
56. Vazirani, V.V.: Approximation Algorithms. Springer, New York (2001)