

Computing Bi-Lipschitz Outlier Embeddings into the Line

Karine Chubarian

Department of Mathematics, Statistics and Computer Science,
University of Illinois at Chicago, IL, USA
kchuba2@uic.edu

Anastasios Sidiropoulos

Department of Computer Science, University of Illinois at Chicago, IL, USA
sidiropo@uic.edu

Abstract

The problem of computing a bi-Lipschitz embedding of a graphical metric into the line with minimum distortion has received a lot of attention. The best-known approximation algorithm computes an embedding with distortion $O(c^2)$, where c denotes the optimal distortion [Bădoiu et al. 2005]. We present a bi-criteria approximation algorithm that extends the above results to the setting of *outliers*.

Specifically, we say that a metric space (X, ρ) admits a (k, c) -embedding if there exists $K \subset X$, with $|K| = k$, such that $(X \setminus K, \rho)$ admits an embedding into the line with distortion at most c . Given $k \geq 0$, and a metric space that admits a (k, c) -embedding, for some $c \geq 1$, our algorithm computes a $(\text{poly}(k, c, \log n), \text{poly}(c))$ -embedding in polynomial time. This is the first algorithmic result for outlier bi-Lipschitz embeddings. Prior to our work, comparable outlier embeddings were known only for the case of additive distortion.

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms

Keywords and phrases metric embeddings, outliers, distortion, approximation algorithms

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2020.36

Category APPROX

Funding *Anastasios Sidiropoulos*: Supported by NSF grants CCF-1815145, CCF-1934915, and by NSF CAREER award 1453472.

1 Introduction

The theory of metric embeddings provides an extensive toolbox that has found applications in several geometric data-analytic tasks. At the high level, an embedding of a metric space $\mathcal{M} = (X, \rho)$ into some metric space $\mathcal{M}' = (X', \rho')$ is a mapping $f : X \rightarrow X'$ that preserves certain interesting geometric properties of \mathcal{M} . In most cases, it is desirable to obtain embeddings that minimize some notion of *distortion*.

Despite the success of metric embeddings methods in several application domains, one significant limitation of most existing methods is that they are not robust to noise in the form of outlier points in the input. This setting is of particular interest in the case where the data does not perfectly fit the underlying geometric model, or when some points are corrupted due to measurement errors. The outlier model also has connections to the setting of adversarial machine learning [13]. More specifically, in the setting of *poisoning attacks*, it is often assumed that a small subset of the training data set is corrupted adversarially. For example, in a classification application, some of the training samples can be modified arbitrarily. Therefore, it is important to design data-analytic primitives that are robust against this type of adversarial input perturbation.

Our aim is to bypass the limitations of current metric embedding methods by designing approximation algorithms that given some input space \mathcal{M} , they compute a small subset of points to delete, and an embedding of the residual space into some desired host space.



© Karine Chubarian and Anastasios Sidiropoulos;
licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2020).

Editors: Jarosław Byrka and Raghu Meka; Article No. 36; pp. 36:1–36:21



Leibniz International Proceedings in Informatics

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

1.1 Our contribution

We now formally define outlier embeddings and state our main result. Let $\mathcal{M} = (X, \rho)$, $\mathcal{M}' = (X', \rho')$ be metric spaces. An injection $f : X \rightarrow X'$ is called an *embedding*. Given an embedding f , its *distortion* is defined as

$$\text{distortion}(f) = \sup_{x \neq y \in X} \frac{\rho'(f(x), f(y))}{\rho(x, y)} \cdot \sup_{x' \neq y' \in X} \frac{\rho(x', y')}{\rho'(f(x'), f(y'))}.$$

We also refer to this notion of distortion as *multiplicative distortion*. An embedding is *bi-Lipschitz* if its distortion is bounded. When $\mathcal{M}' = (\mathbb{R}, \ell_2)$ then we say that \mathcal{M} admits an embedding into the line. If $\text{distortion}(f) \leq c$, then we say that f is a *c-embedding*. We use the following definition for outlier embeddings (see also [19]).

A metric space $\mathcal{M} = (X, \rho)$ admits a (k, c) -embedding into another metric space $\mathcal{M}' = (X', \rho')$ for some $c \geq 1$, $k \geq 0$ if there exists $K \subseteq X$, with $|K| \leq k$, and $f : X \setminus K \rightarrow X'$, with $\text{distortion}(f) \leq c$. We say that such $K \subseteq X$ is an *outlier set* (w.r.t. f).

In the present work, we focus on the case where the input metric space is the shortest-path metric of an unweighted graph, and the host space is the real line. This setting, but without outliers, has been studied extensively in the literature (see Section 1.2 for a more detailed discussion). The shortest-path metrics of unweighted graphs arise naturally in applications, for example, when considering the k -NN graph of a point set; that is, by taking the set of vertices to be a set of samples from some unknown manifold, and the edge set to be all pairs $\{u, v\}$, where u is one of the k nearest neighbors of v . Moreover, the case of embedding into the real line is a prototypical mathematical model for the problem of discovering 1-dimensional structure in a metrical data set.

The following summarizes the main result of this paper.

► **Theorem 1.** *Let G be a graph, $k \geq 0$, $c \geq 1$. There exists a polynomial-time algorithm which given G , k , and c , terminates with exactly one of the following outcomes:*

- (1) *Correctly decides that G does not admit a (k, c) -embedding into the line.*
- (2) *Computes a $(O(c^6 k \log^{5/2} n), O(c^{13}))$ -embedding of G into the line.*

1.2 Related work

Low-distortion metric embeddings have been studied extensively within mathematics and computer science. We refer the reader to [14] for a detailed exposition of the work that is of main interest for computer science. Here, we discuss some results relevant to our work.

Approximation algorithms. The problem of computing an embedding of some input metric space \mathcal{M} into some host space \mathcal{M}' with approximately minimum distortion has received a lot of attention. Most positive results are concerned with the case where \mathcal{M}' is the line, or, more generally, some 1-dimensional space. Specifically, Bădoiu et al. [6] obtained an algorithm which given an unweighted graph that admits a c -embedding into the line, computes a $O(c^3)$ -embedding into the line. Approximation algorithms have also been obtained by Bădoiu et al. [5] for the case where the input is a weighted tree, and by Nayyeri and Raichel [17] for the case where the input is a general metric space.

Approximation algorithms for embedding into more general 1-dimensional spaces have also been considered. Bădoiu et al. [3] consider the case where the host space is a tree, Chepoi et al. [8] consider the case where the host space is an outerplanar graph, and Nayyeri and Raichel [18] generalize this to the case where the host space is a graph of bounded treewidth. Carpenter et al. [7] obtain an approximation algorithm for embedding unweighted graphs into subdivisions of any fixed “pattern” graph H (embedding into the line corresponds to H being a single edge, while embedding into a cycle is the case where H is a triangle).

The case of higher-dimensional host spaces appears to be significantly more challenging. The only positive results are an approximation algorithm for embedding finite subsets of the 2-sphere into \mathbb{R}^2 [6], and approximation algorithms for embedding ultrametrics into \mathbb{R}^d [4, 10]. On the negative side, it is shown that for any $d \geq 1$, the problem of embedding into d -dimensional Euclidean space with minimum distortion is hard to approximate within a factor of $n^{\alpha/d}$, for some constant $\alpha > 0$ (the case $d = 1$ is due to [5] and $d \geq 2$ is due to [16]).

FPT algorithms. The problem of computing an embedding into the line parameterized by the optimal distortion has also been considered. Fellows et al. [12] gave an FPT algorithm for embedding unweighted graphs into the line. A nearly-matching lower bound on the running time (assuming ETH) was obtained by Lokshtanov et al. [15]. FPT algorithms for embedding unweighted graphs into subdivisions of an arbitrary fixed pattern graph H have also been obtained by Carpenter et al. [7].

Outlier embeddings. The problem of computing outlier embeddings was introduced by Sidiropoulos et al. [19]. They considered the case of embedding into d -dimensional Euclidean space, and into trees. The main difference with our work is that [19] deals with the case of *additive* distortion, while we are concerned with multiplicative distortion. As a result, the results in [19] are incomparable to ours. We remark, however, that the case of multiplicative distortion is known to be significantly more challenging. To the best of our knowledge, our result is the first non-trivial upper bound for computing outlier embeddings minimizing the multiplicative distortion.

1.3 High-level overview of the algorithm

We now give an informal description of our algorithm, highlighting the main technical challenges. The input consists of an undirected graph G and some $k \geq 0$, $c \geq 1$. The algorithm either correctly decides that there exists no (k, c) -embedding of G into the line, or outputs a (k', c') -embedding of G into the line, for some $k' = \text{poly}(k, c, \log n)$, $c' = \text{poly}(c)$.

The crux of the algorithm is to identify and remove three “obstructions” for low-distortion embeddability into the line. These three obstructions are *regions of high density*, *large metrical cycles* and *large metrical tripods*. We next discuss the steps used to handle each one of these obstructions, and describe how all the steps are combined in the final algorithm.

Obstruction 1: Reducing the density. The *density* of a graph is defined to be


$$\Delta(G) = \max_{v \in V(G), R \in \mathbb{N}} \frac{|\text{Ball}_G(v, R)| - 1}{2R}.$$

It is known that the density of any graph that admits a c -embedding into the line is $O(c)$ [6]. Therefore, if G admits a (k, c) -embedding, then there must exist some set of at most k vertices, whose deletion leaves a graph with density $O(c)$. We observe that the density of a graph is a hereditary property, meaning that for any $H \subseteq G$, we have $\Delta(H) \leq \Delta(G)$. This leads to a following recursive procedure: if the density is higher than $O(c)$, we compute a balanced vertex separator $X \subseteq V(G)$, and recurse on $G \setminus X$. We set

$$K_{\text{density}} := \bigcup_{\text{all separators } X} X.$$

Let us also denote $G \setminus K_{\text{density}}$ as G' . It is immediate that $\Delta(G') = O(c)$, and we show that $|K_{\text{density}}| = \text{poly}(k, c, \log n)$.

Obstruction 2: Eliminating large metrical cycles. It is known that any embedding of the n -cycle into the line must incur distortion $\Omega(n)$ [6]. More generally, it is possible to define an obstruction, which we refer to as a *metrical cycle*, and which contains cycles as a special case, but allows for more general shortest-path distances (see Figure 1). We show how to delete a small number of vertices so that the resulting graph does not contain any large metrical cycles, and then we find a low-distortion embedding into some forest.




■ **Figure 1** Example of a large metrical cycle.

We now briefly describe the procedure for eliminating large metrical cycles. We start by computing a $\text{poly}(c)$ -net N in G' . We then find a Voronoi partition \mathcal{P} centered at N : for any vertex $v \in G'$, we assign v to a cluster centered at its nearest neighbour $y \in N$ (we break ties to ensure connectivity). Let H be the minor of G obtained by contracting each cluster to its center $y \in N$. We compute an approximate minimum feedback vertex set Y in H . We set

$$K_{\text{forest}} := \bigcup_{x \in Y} \mathcal{P}(x),$$

and $G'' = G' \setminus K_{\text{forest}}$. Note that the low density of G' ensures that $|K_{\text{forest}}|$ is small. Furthermore, we show that G'' admits a low-distortion embedding into a forest.




■ **Figure 2** Elimination of large metrical cycles. From left to right: the graph G' , the minor H , the forest $H \setminus Y$, and the graph G'' .

Obstruction 3: Eliminating large metrical tripods. A *tripod* is a tree consisting of the union of three paths with a common endpoint; we say that a tripod is R -large if the length of each of the three paths is at least R . Any embedding of a R -large tripod into the line must incur distortion $\Omega(R)$. We show how to delete a small number of vertices so that the resulting graph does not have any subgraphs with a shortest-path metric that resembles that of a $\Omega(\text{poly}(c))$ -large tripod. More specifically, via a reduction to the Minimum Set Cover problem, we compute some $Z \subseteq V(H \setminus Y)$, so that the forest $H \setminus (Y \cup Z)$ does not contain any $\Omega(\text{poly}(c))$ -large tripods (see Figure 3). We set

$$K_{\text{tripod}} := \bigcup_{w \in (H \setminus Y) \setminus Z} \mathcal{P}(w).$$

and $G'' = G' \setminus K_{\text{tripod}}$. Since the forest $H \setminus (Y \cup Z)$ does not contain any large tripods, we can show that it admits a low-distortion embedding into the line. Furthermore, we can use this embedding to also embed G'' into the line.



■ **Figure 3** Elimination of a large tripod. A yellow vertex removes the red tripod and the yellow dotted tripod simultaneously.

Putting everything together. The final algorithm combines the above procedures for eliminating the three obstructions that we have identified. At each obstruction elimination step, we remove a small set of vertices. One additional complication is that, because c -embeddability into the line is not a hereditary property, this can produce a graph that does not admit a low-distortion embedding into the line. We show that this issue can be avoided by deleting a slightly larger superset of vertices, which eliminates the obstruction at hand, while maintaining the existence of a low-distortion embedding.

1.4 Organization

The rest of the paper is organized as follows. We introduce necessary notation and definitions in Section 2. In Section 3, we present our main algorithm and we state the main technical results needed. In Section 3.2 we prove a technical lemma which will be applied throughout the paper. Sections A, B, C, D elaborate on the subroutines executed by the main algorithm.

2 Preliminaries

2.1 Graphs

Given a graph G , we refer to its vertex set as $V(G)$ and to its edge set as $E(G)$. For any $C \subseteq V(G)$, we denote by $G[C]$ the subgraph of G induced on C . Let d_G denote the shortest-path distance of G ; unless otherwise noted, we assume that all edges in G are undirected and have unit length.


► **Definition 2** (Local density). For any $v \in V(G)$ and $R \in \mathbb{N}$, we define

$$\Delta_G(v, R) = \frac{|\text{Ball}_G(v, R)| - 1}{2R}$$

The local density of the graph G is defined to be

$$\Delta(G) = \max_{v \in V(G), R \in \mathbb{N}} \Delta_G(v, R).$$

► **Definition 3** (Tripod). Let G be a graph, $R \geq 1$, $v, v_1, v_2, v_3 \in V(G)$, and let P_1, P_2, P_3 be paths in G , where for all $i \in [3]$, P_i is a path with endpoints v and v_i . Suppose that for all $i \neq j \in [3]$, and for all $u \in P_j$, we have $d_G(v_i, u) \geq R$. In other words, each endpoint v_i is at distance at least R from every vertex in the other two paths. Then we say that the tree $P_1 \cup P_2 \cup P_3$ is a R -tripod with root v (in G).



■ **Figure 4** A tripod rooted at v with leaves v_1, v_2, v_3 .

2.2 Some useful approximation results

For a graph G , a *feedback vertex set* is some $X \subseteq V(G)$, such that $G \setminus X$ is acyclic. In the Minimum Feedback Vertex Set problem we are given a graph G and the goal is to find a feedback vertex set in G of minimum cardinality. We recall the following result on approximating the Minimum Feedback Vertex Set problem.

► **Theorem 4** (Bafna et al. [1]). *There exists a polynomial-time 2-approximation algorithm for the Minimum Feedback Vertex Set problem.*

Given a graph G and some $\alpha \in [0, 1]$, we say that some $X \subseteq V(G)$ is a α -balanced vertex separator (of G) if every connected component of $G \setminus X$ has at most $\alpha \cdot |V(G)|$ vertices. We recall the following algorithmic result on computing balanced vertex separators.

► **Theorem 5** (Feige et al. [11]). *There exists a polynomial-time algorithm which given a graph that admits a $2/3$ -balanced vertex separator of size s , outputs a $3/4$ -balanced vertex separator of size at most $O(\sqrt{\log n} \cdot s)$.*

Recall that an instance to the Minimum Set Cover problem consists of some set U (the universe), and a set \mathcal{C} of subsets of U . The goal is to find a subset of \mathcal{C} of minimum cardinality that covers U .

► **Theorem 6** (Chvátal [9]). *There exists a polynomial-time $O(\log n)$ -approximation algorithm for the Minimum Set Cover problem.*

2.3 Voronoi minors

For some metric space $\mathcal{M} = (X, \rho)$, and some $R > 0$, we say that some $N \subseteq X$ is a R -net of M if for any $p, q \in N$, $\rho(p, q) > R$, and $X \subseteq \bigcup_{p \in N} \text{Ball}_{\mathcal{M}}(p, R)$. For a graph G , we say that some $N \subseteq V(G)$ is a R -net of G if N is a R -net of the shortest-path metric of G .


► **Definition 7** (Graphical Voronoi partition). *Let G be a graph, and let $Y \subseteq V(G)$. Let \mathcal{P} be a partition of $V(G)$ satisfying the following conditions:*

- (1) *Every cluster in \mathcal{P} contains exactly one vertex in Y .*
- (2) *For any $v \in V(G)$, the cluster containing v , $\mathcal{P}(v)$, also contains some nearest neighbor of v in Y .*
- (3) *For any cluster $C \in \mathcal{P}$, we have that $G[C]$ is connected.*

We say that \mathcal{P} is a Voronoi partition of G centered at Y .

We note the following easy fact.

► **Lemma 8.** *For any graph G , and $Y \subseteq V(G)$, there exists a Voronoi partition \mathcal{P} of G centered at Y .*



■ **Figure 5** A Voronoi partition centred at 3-net $N = \{y_1, y_2, y_3\}$ and a corresponding 3-minor.

Proof. Construct \mathcal{P} by assigning each $v \in V(G)$ to the cluster containing its nearest neighbor in Y . In order to ensure that each cluster C induces connected subgraph $G[C]$ it suffices to ensure that shortest-paths in G are unique. This can be achieved by breaking ties between different paths lexicographically (viewing paths as sequences of vertices with unique integer labels) (see also [6]). ◀

► **Definition 9** (*R*-Minor). Let G be a graph, $R > 0$, and let N be a R -net of G . Let \mathcal{P} be a Voronoi partition of G centered at N . Let H be the minor of G obtained by contracting each cluster in \mathcal{P} into the unique net point in C . Then we say that \mathcal{P} is a R -partition and H is a R -minor of G induced by \mathcal{P} (see Figure 5 for an example).

3 The Main Algorithm

In this section we present and analyze the main algorithm of the paper. For the clarity, we first state some key technical ingredients used by the algorithm. We then present the main algorithm and its analysis. The proofs of the technical ingredients appear in latter Sections.

3.1 Technical ingredients used by the main algorithm

Density reduction. The first technical ingredient used by the main algorithm is a procedure for reducing the local density of the input graph. This is summarized in Lemma 10. Its proof is given in Section A.


► **Lemma 10** (Density Reduction). *There exists a polynomial-time algorithm given a graph G , $k \geq 0$, $c \geq 1$, terminates with exactly one of the following outcomes:*

- (1) *Correctly decides that G does not admit a (k, c) -embedding into the line.*
- (2) *Outputs some $Y \subseteq V(G)$ such that $\Delta(G \setminus Y) \leq c$, with $|Y| = O(ck \log^{3/2} n)$. In particular, if $\Delta(G) \leq c$, then the algorithm outputs \emptyset .*

Eliminating large metrical cycles. The next technical ingredient is a procedure for eliminating large metrical cycles. This is summarized in Lemma 11, whose proof is given in Section B.

► **Lemma 11** (Embedding into a forest). *There exists a polynomial-time algorithm which given a graph G , $c \geq 1$, and $k \geq 0$, terminates with exactly one of the following outcomes:*

- (1) *Correctly decides that G does not admit a (k, c) -embedding into the line.*
- (2) *Outputs a c -net N of G , a c -partition \mathcal{P} centered at N , a c -minor H induced by \mathcal{P} , and some feedback vertex set X of H , with $|X| \leq 2k$.*



■ **Figure 6** A $3 \times n$ grid G can be embedded into the line with distortion $O(1)$; one could follow the red dotted path on the grid and embed the vertices consecutively. A yellow line depicts U . Now, if we delete a yellow vertex from $G \setminus U$, the resulting graph will be just a path.

Eliminating large metrical tripods. The next obstruction that the main algorithm needs to remove is large metrical tripods. This is done using Lemmas 12 and 13. Their proofs appear in Section C.

► **Lemma 12** (Tripods as obstructions to embeddability). *Let G be a graph, $R \geq 1$, and let J be a R -tripod in G . Then for any c -embedding of G into the line we have $c \geq 2R$.*

► **Lemma 13** (Tripod elimination). *There exists a polynomial-time algorithm which given a forest F , $R \geq 1$, $k \geq 0$, terminates with exactly one of the following outcomes:*

- (1) *Correctly decides that there exists no $X' \subseteq V(F)$, with $|X'| \leq k$, such that $F \setminus X'$ does not contain any R -tripod as a subgraph.*
- (2) *Outputs some $X' \subseteq V(F)$, with $|X'| = O(k \log n)$, such that $F \setminus X'$ does not contain any R -tripod as a subgraph.*

Embedding a tree with no large tripods into the line. Once all the obstructions have been removed, the problem is reduced to computing an embedding of a tree with no large tripods into the line. This is done using Lemma 14, whose proof appears in Section D.

► **Lemma 14.** *Let $R \geq 1$, and let T be a tree that does not contain any R -tripod as a subgraph. Then T admits a $O(\Delta(T) \cdot R)$ -embedding into the line. Moreover, this embedding can be computed in polynomial time.*

3.2 The Repairing Lemma

The main algorithm proceeds in several steps. At each step, it uses some of the procedures described above to delete small subsets of vertices. However, because c -embeddability into the line is not a hereditary property, it is possible that the deletion of some small set of vertices destroys some candidate solution. As an illustrative example, let G be the $3 \times (n/3)$ grid. Note that G admits a $O(1)$ -embedding into the line (i.e. without outliers). This embedding can be realized by consecutively traversing the columns of the grid. Let U be the set of vertices that do not lie on the outer boundary cycle of G . Then, $G \setminus U$ is the $(2n/3 + 2)$ -cycle, and therefore any embedding of $G \setminus U$ into the line has distortion $\Omega(n)$. However, by removing one additional vertex from $G \setminus U$ we obtain a path, which admits a 1-embedding into the line (see Figure 6). We show that the above “repairing” process can be performed for arbitrary U . First, we prove two auxiliary statements.

► **Lemma 15.** *Let G be a graph, $k > 0$, $c > 1$. Assume that G admits a (k, c) -embedding into a line. Suppose G admits a (k, c) -embedding into the line realized by $f : G \setminus K \rightarrow \mathbb{R}$. Then, there exists a (k, c) -embedding f' of G into a line such that if $j > i$ then for any $v \in G_i, w \in G_j$ we have $f'(w) > f'(v)$.*

Proof. Let

$$v_1 = \arg \min_{v \in V(G) \setminus K} \{f(v)\}$$

$$v_2 = \arg \max_{v \in V(G) \setminus K} \{f(v)\}$$

and let $M = f(v_2) - f(v_1)$. Without loss of generality, we can assume that $f(v_1) = 0$ and $f(v_2) = M$ by setting $f(v) := f(v) - f(v_1)$. For each $v \in G_i$ we define $f'(v) = f(v) + 2i \cdot M$. We claim that f' and f have the same distortion. If $v, w \in G_i$ then we have

$$|f'(w) - f'(v)| = |(f(w) + 2i \cdot M) - (f(v) + 2i \cdot M)| = |f(w) - f(v)|.$$

If $v \in G_i$ and $w \in G_j$ for $i \neq j$ then the distance between them in the embedding does not contribute to the distortion.

It remains to show that $f'(w) > f'(v)$ for all $w \in G_j, v \in G_i$ with $j > i$. We have

$$f'(w) - f'(v) = f(w) - f(v) + 2(j - i)M > -M + 2M > 0$$

and the claim follows by induction. \blacktriangleleft

Let $G = (V(G), E(G))$ be a graph and let $f : G \rightarrow \mathbb{R}$. Consider $Z = \{v_1, \dots, v_m\} \subseteq V(G)$ such that $f(v_1) < f(v_2) < \dots < f(v_m)$. Then Z is *consecutive with respect to f* if for all $w \in V(G) \setminus Z$ either $f(w) < f(v_1)$ or $f(v_m) < f(w)$.

► Lemma 16. *Let G be a graph, $c > 0$. Assume that G admits a $(0, c)$ -embedding into the line realized by $f : G \rightarrow \mathbb{R}$. Let $Z = \{z_1, \dots, z_m\} \subseteq V$ be consecutive with respect to f . Suppose that $f(v_m) - f(v_1) \geq c$; then Z is a vertex separator in G .*

Proof. We claim that

$$X = \{x \in V(G) \mid f(x) < f(v_1)\}$$

$$Y = \{y \in V(G) \mid f(v_k) < f(y)\}$$

are disconnected in $G \setminus Z$. Assume otherwise; then there exists $\{x, y\} \in E(G)$ with $x \in X, y \in Y$. Thus

$$|f(y) - f(x)| = f(y) - f(v_k) + f(v_k) - f(v_1) + f(v_1) - f(x) \geq c + 1 > c \cdot d_G(x, y)$$

which contradicts the distortion assumption. \blacktriangleleft

We can now prove the Repairing Lemma:

► Lemma 17 (Repairing Lemma). *Let G be a graph, $U \subset V(G)$, $k \geq 0, c \geq 1$. Suppose that G admits a (k, c) -embedding into a line. Then, $G \setminus U$ admits a $((2c+1)|U| + k, 4c^3 + c)$ -embedding into a line.*

Proof. Let f be a (K, c) -embedding of G into the line, with $|K| = k$. Let $U' = U \cap K$, and $U'' = U \setminus K$. For any $v \in U \setminus K$, let


$$I_{\text{inner}}(v) = \text{Ball}_{\mathbb{R}}(f(v), c),$$

$$I_{\text{outer}}(v) = \text{Ball}_{\mathbb{R}}(f(v), 2c^2) \setminus I_{\text{inner}}(v),$$

$$V_{\text{inner}}(v) = \{u \in V(G) \setminus K : f(u) \in I_{\text{inner}}(v)\}$$

$$V_{\text{outer}}(v) = \{u \in V(G) \setminus K : f(u) \in I_{\text{outer}}(v)\}$$

36:10 Computing Bi-Lipschitz Outlier Embeddings into the Line



■ **Figure 7** Inner, outer, safe and exposed vertices with respect to v for $c = 2$.

Let also

$$\begin{aligned}
 I_{\text{inner}} &= \bigcup_{v \in U \setminus K} I_{\text{inner}}(v) \\
 I_{\text{outer}} &= \left(\bigcup_{v \in U \setminus K} I_{\text{outer}}(v) \right) \setminus I_{\text{inner}} \\
 V_{\text{inner}} &= \{u \in V(G) \setminus K : f(u) \in I_{\text{inner}}\}, \\
 V_{\text{outer}} &= \{u \in V(G) \setminus K : f(u) \in I_{\text{outer}}\}, \\
 V_{\text{exposed}} &= V_{\text{inner}} \cup V_{\text{outer}}, \\
 V_{\text{safe}} &= V(G) \setminus V_{\text{exposed}}.
 \end{aligned}$$

We can now define

$$K' = K \cup V_{\text{inner}}.$$

Since the minimum distance in G is one, and f is non-contracting, it follows that

$$|K'| \leq |K| + (2c + 1)|U|.$$

Let $c' = (4c^3 + c)$. It remains to construct any (K', c') -embedding f' . By lemma 15 it is enough to construct a c' -embedding for each connected component of $G \setminus K'$

We may thus focus on any connected component C of $G \setminus K'$. Let $f' = (4c^2 + 1) \cdot f|_C$ (that is, f' is the restriction of f on C scaled by a factor of $4c^2 + 1$). It suffices to show that f' is a $(4c^3 + c)$ -embedding of C .

If there exist $v \in U \setminus K$, and $u \in C$ such that $f(v) < f(u)$, then we set

$$z_L = \arg \max_{v \in U \setminus K : \forall u \in C, f(z_L) < f(u)} \{f(v)\},$$

Similarly, if there exist $v \in U \setminus K$, and $u \in C$ such that $f(v) > f(u)$, then we set

$$z_R = \arg \min_{v \in U \setminus K : \forall u \in C, f(z_R) > f(u)} \{f(v)\}.$$

Let $u, v \in V(C)$. We first bound the expansion of f' . Since $K \subset K'$, it follows what $d_{G \setminus K}(u, v) \leq d_{G \setminus K'}(u, v)$, and thus

$$\begin{aligned}
 |f'(u) - f'(v)| &= (4c^2 + 1) \cdot |f(u) - f(v)| \leq (4c^3 + c) \cdot d_{G \setminus K}(u, v) \\
 &\leq (4c^3 + c) \cdot d_{G \setminus K'}(u, v).
 \end{aligned} \tag{1}$$

It remains to show that f' is non-contractive. Let P be the shortest path between u and v in $G \setminus K$. Let us first assume that $u, v \in V_{\text{safe}}$; we will consider the general case later. If z_L is defined and $P \cap V_{\text{outer}}(z_L)$, we first construct a new path P' that avoids $V_{\text{outer}}(z_L)$, as

follows. When traversing P starting from u , let u_1 be the last vertex before visiting $V_{\text{outer}}(z_L)$ for the first time; let also u_2 be the first vertex visited immediately after leaving $V_{\text{outer}}(z_L)$ for the last time.

Since the expansion of f is at most c , it follows that

$$\begin{aligned} f(u_1) &\in (f(z_L) + 2c^2, f(z_L) + 2c^2 + c], \\ f(u_2) &\in (f(z_L) + 2c^2, f(z_L) + 2c^2 + c], \end{aligned}$$

and thus

$$d_{G \setminus K}(u_1, u_2) \leq |f(u_1) - f(u_2)| \leq c. \quad (2)$$

Let W be the shortest path between u_1 and u_2 in $G \setminus K$. Since every edge of W is stretched by at most a factor of c in f , it follows by (2) that W cannot enter $V_{\text{inner}}(z_L)$, and thus $W \subseteq G \setminus K'$. Therefore $d_{G \setminus K'}(u_1, u_2) = d_{G \setminus K}(u_1, u_2) \leq c$. We can replace P by the path $P' := P[u, u_1] \circ W \circ P[u_2, v]$, which does not intersect $v_{\text{outer}}(z_L)$. We obtain that $\text{length}(P') = \text{length}(P[u, u_1]) + \text{length}(W) + \text{length}(P[u_2, v]) \leq c + \text{length}(P) \leq c + d_{G \setminus K}(u, v)$.

Next, if z_R exists and $P' \cap V_{\text{outer}}(z_R) \neq \emptyset$, then via a symmetric process we can replace P' by a new path P'' between u and v in $G \setminus K$ avoids $V_{\text{outer}}(z_R) \cup V_{\text{outer}}(z_L)$, with

$$\text{length}(P'') \leq \text{length}(P') + c \leq \text{length}(P) + 2c.$$

This implies that $P' \subseteq G \setminus K'$.

We therefore obtain

$$\begin{aligned} |f'(u) - f'(v)| &= (4c^2 + 1) \cdot |f(u) - f(v)| \geq (4c^2 + 1) \cdot d_{G \setminus K}(u, v) \\ &\geq (4c^2 + 1) \cdot (d_{G \setminus K'}(u, v) - 2c) > d_{G \setminus K'}(u, v). \end{aligned} \quad (3)$$

By (1) and (3) we obtain that f is a $(4c^3 + c)$ -embedding of $G \setminus K'$, as required.

It remains to consider the case where either $u \in V_{\text{exposed}}$, or $v \in V_{\text{exposed}}$. Let Q be a shortest path between u and v in $G \setminus K'$. If $Q \cap V_{\text{safe}} = \emptyset$, then $\text{length}(Q) \leq 4c^2 - 2c$, thus

$$d_{G \setminus K'}(u, v) \leq (4c^2 - 2c) \leq (4c^2 - 2c) \cdot d_{G \setminus K'}(u, v),$$

which implies that f' is non-contractive, as required. We may therefore assume for the remainder of the proof that $Q \cap V_{\text{safe}} \neq \emptyset$. If $u \in V_{\text{exposed}}$, then we may assume w.l.o.g. that $u \in V_{\text{outer}}(z_L)$. When traversing P starting from u , let u_1 be the first vertex visited immediately after leaving $V_{\text{outer}}(z_L)$. When traversing Q starting from u , let u_2 be the first vertex visited in V_{safe} . By an argument identical to the one used in the previous case, we can obtain a new path between u and v , given by $Q[u, u_2] \circ W \circ P[u_1, v]$, where $\text{length}(Q[u, u_2]) \leq 2c^2 - c$ (since all vertices in $Q[u, u_2]$ except the last one are contained in the rightmost segment of $V_{\text{outer}}(z_L)$), $W \subseteq G \setminus K'$, and $\text{length}(W) \leq c$ (as in the previous case). We thus obtain a path of length at most $d_{G \setminus K}(u, v) + 2c^2$. If $v \in V_{\text{exposed}}$, we repeat the above process after exchanging u and v . We thus arrive at a path between u and v of length at most $d_{G \setminus K}(u, v) + 4c^2 \leq (4c^2 + 1) \cdot d_{G \setminus K}$, which does not intersect V_{inner} , and thus it is contained in $G \setminus K'$. It follows that f' is non-contractive, and thus a $(4c^3 + c)$ -embedding, which concludes the proof. \blacktriangleleft

3.3 The algorithm

Given the technical ingredients presented above, we are now ready to describe our main algorithm. Recall that the input consists of a graph G , and $k \geq 0$, $c \geq 1$. The algorithm proceeds in the following steps.

36:12 Computing Bi-Lipschitz Outlier Embeddings into the Line

Step 1: Density reduction. Using the algorithm from Lemma 10 we can either correctly decide that G does not admit a (k, c) -embedding into the line, in which case we terminate, or we compute some $X_{\text{density}} \subseteq V(G)$, with $|X_{\text{density}}| \leq O(ck \log^{3/2} n)$, such that $\Delta(G \setminus X_{\text{density}}) \leq c$.

Step 2: Cycle elimination. Let $k' = (2c + 1)|X_{\text{density}}| + k$ and $c' = 4c^3 + c$. Using the algorithm from Lemma 11 we either correctly decide that G' does not admit a (k', c') -embedding into the line, or we compute a c' -net N of G' , a c' -partition \mathcal{P} centered at N , a c' -minor H induced by \mathcal{P} , and some feedback vertex set Y_{forest} of H , with $|Y_{\text{forest}}| \leq 2k'$. If G' does not admit a (k', c') -embedding into the line, then we terminate by deciding that G does not admit a (k, c) -embedding into the line.

Step 3: Tripod elimination. Let $F = H \setminus Y_{\text{forest}}$, and recall that Y_{forest} is a feedback vertex set for H , and thus F is a forest. Using the algorithm from Lemma 13, in polynomial time, we either decide that there exists no $Y_{\text{tripod}} \subseteq V(F)$, with $|Y_{\text{tripod}}| \leq k'$, such that $F \setminus Y_{\text{tripod}}$ does not contain any $(c'/2 + 1)$ -tripod, in which case we terminate deciding that G does not admit a (k, c) -embedding into the line, or we compute some $Y_{\text{tripod}} \subseteq V(F)$, with $|Y_{\text{tripod}}| = O(k \log n)$, such that $F \setminus Y_{\text{tripod}}$ does not contain any $(c'/2 + 1)$ -tripods.

Step 4: Embedding into a forest. Let $F' = F \setminus Y_{\text{tripod}}$. Let

$$X_{\text{forest}} = \bigcup_{v \in Y_{\text{forest}}} \mathcal{P}(v),$$

$$X_{\text{tripod}} = \bigcup_{v \in Y_{\text{tripod}}} \mathcal{P}(v),$$

and

$$K = X_{\text{density}} \cup X_{\text{forest}} \cup X_{\text{tripod}}.$$

Let F'' be the forest obtained from F' as follows. Initially, we set $F'' := F'$. For each $v \in V(G) \setminus K$, let $u(v)$ be the unique vertex in $N \cap \mathcal{P}(v)$; we add v to F'' as a leaf attached to $u(v)$. This completes the construction of the forest F'' .

Step 5: Embedding into the line. Finally, we compute an embedding f of F'' into the line using the algorithm from Theorem 14. We output the embedding $\varphi := 2c'c \cdot f$ (that is, f scaled by a factor of $2c'c$).

3.4 Analysis of the main algorithm

We now analyze the main algorithm presented above. First, we state some auxiliary properties of c -minors and c -partitions.

► **Lemma 18.** *Let G be a graph, $R \geq 1$. Let N be a R -net of G , \mathcal{P} a corresponding R -partition and H a R -minor G induced by \mathcal{P} . Then for any $Y \subseteq V(H)$ all of the following hold:*

- (1) $N' := N \setminus Y$ is a R -net in $G' := G \setminus (\cup_{v \in Y} \mathcal{P}(v))$
- (2) $\mathcal{P}' := \mathcal{P} \setminus (\cup_{v \in Y} \mathcal{P}(v))$ is the R -partition of G' centered at N'
- (3) $H' := H \setminus Y$ is the R -minor of G' induced by \mathcal{P}' .

Proof. We first show (1). Since by deleting vertices the shortest-path distances cannot increase, we have that for all $u, v \in N'$, $d_{G'}(u, v) \geq d_G(u, v) > R$. It thus remains to show that for any $x \in V(G')$ there exists $v \in N'$ such that $d_{G'}(x, v) \leq R$. Consider an arbitrary

$x \in V(G')$. Let $v \in N$ be such that $x \in \mathcal{P}(v)$. Since the shortest path between v and x in G is contained in $\mathcal{P}(v)$, it follows that

$$d_{G'}(x, v) \leq d_{G'[\mathcal{P}(v)]}(x, v) = d_{G[\mathcal{P}(v)]}(x, v) = d_G(x, v) \leq c,$$

which implies that N' is a c -net of G' .

Next, we show (2). Since for all $v \in N'$, we have $\mathcal{P}'(v) = \mathcal{P}(v)$, it follows that \mathcal{P}' is a partition of $V(G')$. Since by (1) N' is a R -net of G' , and for all $v \in N'$, and for all $x \in \mathcal{P}'(v)$ we have $d_{G'}(v, x) \leq R$, it follows that \mathcal{P}' is a R -partition of G' centered at N' .

Finally, we show (3). Let \tilde{H} be the R -minor of G' induced by \mathcal{P}' . We prove that $V(H') = V(\tilde{H})$ and $E(H') = E(\tilde{H})$. For the first equality, observe that

$$V(H') = V(H \setminus Y) = N \setminus Y = N' = V(\tilde{H}).$$

It remains to show that $E(H') = E(\tilde{H})$. Consider an arbitrary $\{u, v\} \in E(H')$. Since $H' \subseteq H$ we have that $\{u, v\} \in E(H)$. Then there must exist a path $P \subseteq G$ between u, v with $P \subseteq \mathcal{P}(u) \cup \mathcal{P}(v)$. Since $u, v \in V(H \setminus Y) = N'$ we have that $\mathcal{P}(u) = \mathcal{P}'(u)$, $\mathcal{P}(v) = \mathcal{P}'(v)$. Thus, $P \subseteq \mathcal{P}'(u) \cup \mathcal{P}'(v)$ which yields $\{u, v\} \in E(\tilde{H})$. Now consider an arbitrary $\{u, v\} \in E(\tilde{H})$; it induces a path $Q \subseteq G'$ between u, v such that $Q \subseteq \mathcal{P}'(u) \cup \mathcal{P}'(v)$. Since $\mathcal{P}'(u) = \mathcal{P}(u)$, $\mathcal{P}'(v) = \mathcal{P}(v)$ we obtain $\{u, v\} \in E(H)$. Then from $u, v \in N' = N \setminus Y$ we have $\{u, v\} \in E(H \setminus Y) = E(H')$ which concludes the proof. \blacktriangleleft

► **Lemma 19.** *Let G be a graph and let $R > 0$. Let N be c -net of G , \mathcal{P} a c -partition centered at N , and H a R -minor induced by \mathcal{P} . Then for any $u, v \in N$ we have $d_H(u, v) \leq d_G(u, v)$.*

Proof. Let $P \subseteq G$ be a shortest path between u, v and let $J := \{w \in N : P \cap \mathcal{P}(w) \neq \emptyset\}$. Let $Q \subseteq H$ be a shortest path between u, v . We claim that

$$\text{length}(Q) \leq |J| - 1 \leq \text{length}(P).$$

Assume for contradiction that $\text{length}(Q) > |J| - 1$. Consider arbitrary $\{x_1, x_2\} \in E(P)$ such that $x_1 \in \mathcal{P}(w_1), x_2 \in \mathcal{P}(w_2)$ for $w_1 \neq w_2$; hence $\{w_1, w_2\} \in E(H)$. Therefore, P induces a walk $W \subseteq H$ such that $v, u \in V(W)$. Hence, there is a path $Q' \subseteq W$ such that $v, u \in V(Q')$; note that $\text{length}(Q) \leq |V(W)| - 1 = |J| - 1$. Thus,

$$\text{length}(Q') \leq |J| - 1 < \text{length}(Q) = d_H(v, u).$$

which gives a contradiction, and concludes the proof. \blacktriangleleft

We now have all the necessary ingredients in place to prove Theorem 1, which is the main result of this paper.

Proof of Theorem 1. We analyze the algorithm presented above. By Lemma 10, if we terminate at Step 1, then we correctly decide that G does not admit a (k, c) -embedding. Otherwise, by Lemma 17, it follows that if G admits a (k, c) -embedding into the line, then $G' = G \setminus X_{\text{density}}$ admits a (k', c') -embedding into the line, with $k' = (2c + 1)|X_{\text{density}}| + k = O(c^2 k \log^{3/2} n)$ and $c' = 4c^3 + c$.

By Lemma 11, if we decide that G' does not admit a (k', c') -embedding into the line, then, by the above discussion, this certifies that G does not admit a (k, c) -embedding into the line; we can thus correctly decide this fact in Step 2.

Suppose that G' admits a (k', c') -embedding into the line. Thus, there exists some $K' \subseteq V(G')$, with $|K'| \leq k'$, such that $G' \setminus K'$ admits a c' -embedding into the line. Let J be the set of all $v \in N$ such that the Voronoi cell of v intersects K' , that is $J = \{v \in N : K' \cap \mathcal{P}(v) \neq \emptyset\}$.

36:14 Computing Bi-Lipschitz Outlier Embeddings into the Line

We claim that $F \setminus J$ does not contain any $(3c'/2 + 1)$ -tripod. For the sake of contradiction, suppose that $F \setminus J$ contains some $(3c'/2 + 1)$ -tripod $T = P_1 \cup P_2 \cup P_3$, where P_1, P_2, P_3 are three paths sharing a root r . For any $i \in [3]$ let z_i be the endpoint of P_i other than r . Then for any $i \in [3]$ there exists a path Q_i in $G' \setminus K'$ between r and z_i . We claim that for all $i \neq j \in [3]$, for all $u \in V(Q_j)$, we have $d_{G' \setminus K'}(z_i, u) \geq c'/2 + 1$. By Lemma 18, $F \setminus J$ is a c' -minor of $G' \setminus K'$ with respect to the Voronoi partition \mathcal{P}_J with $\mathcal{P}(w) = \mathcal{P}_J(w)$ for all $w \in V(F \setminus J)$. Let w' be such that $u \in \mathcal{P}_J(w')$. By Lemma 19 obtain

$$\begin{aligned}
 d_{G' \setminus K'}(z_i, u) &\geq d_{G' \setminus K'}(z_i, w') - d_{G' \setminus K'}(w', u) && \text{(by the triangle inequality)} \\
 &\geq d_{G' \setminus K'}(z_i, w') - c' && \text{(since } u \in \mathcal{P}_J(w')\text{)} \\
 &\geq d_{F \setminus J}(z_i, w) - c' && \text{(by Lemma 19)} \\
 &\geq 3c'/2 + 1 - c' && \text{(since } T \text{ is a } (3c'/2 + 1)\text{-tripod)} \\
 &= c'/2 + 1.
 \end{aligned}$$

Therefore, by Lemma 12 we conclude that $G' \setminus K'$ does not admit a c' -embedding into the line, which is a contradiction. Therefore, we have established that if G' admits a (k, c) -embedding into the line, then there exists some $J \setminus V(F)$, with $|J| \leq k'$, such that $F \setminus J$ does not contain any $(3c'/2 + 1)$ -tripods.

Therefore, in Step 3, if we do not find a set Y_{tripod} of the desired size, then we correctly decide that G does not admit a (k, c) -embedding into the line.

Next consider the case where in Step 3 we compute a set Y_{tripod} of the desired size. Since F' does not contain any $(3c'/2 + 1)$ -tripods, it follows by the construction of F'' , that F'' does not contain any $(3c'/2 + 3)$ -tripods (since every leaf in F becomes the center of a star in F'). Moreover, we have $\Delta(F'') \leq \Delta(F') \cdot O(c' \Delta(G'))$, since every vertex in F'' corresponds to a star that contains the vertices of a Voronoi cell in G' , and every such cell has size at most $O(c' \Delta(G'))$. Thus, by Lemma 14 we compute a c'' -embedding of F'' into the line, where $c'' = O(\Delta(F'')c') = O(\Delta(F')c^3 \Delta(G')) = O(\Delta(F)c^3 \Delta(G)) = O(\Delta(H)c^4)$, since $\Delta(\Gamma_1) \leq \Delta(\Gamma_2)$ for all $\Gamma_1 \subset \Gamma_2$. Moreover we have $\Delta(H) \leq \Delta(G') \cdot O(c' \cdot \Delta(G')) = O(c^5)$, since every vertex in H corresponds to a Voronoi cell consisting of at most $O(c' \cdot \Delta(G'))$ vertices. Therefore $c'' = O(c^9)$, and thus we have obtained a $O(c^9)$ -embedding f of F'' into the line. Note that since $V(F'') = V(G \setminus K)$, it follows that f is also a (κ, σ) -embedding of G into the line, where $\kappa = |K|$, for some $\sigma \geq 1$.

It remains to bound κ and σ . We have

$$\kappa = |X_{\text{density}}| + |X_{\text{forest}}| + |X_{\text{tripod}}|.$$

Since G admits a (k, c) -embedding into the line, it follows from Lemma 10 that

$$|X_{\text{density}}| = O(ck \log^{3/2} n).$$

Moreover, $\Delta(G \setminus X_{\text{density}}) \leq c$, thus for any \tilde{c} -partition \mathcal{P} induced by an arbitrary \tilde{c} -net N of $G \setminus X_{\text{density}}$, and any $v \in N$, we have

$$|\mathcal{P}(v)| = O(\tilde{c} \cdot \Delta(G \setminus X_{\text{density}})) = O(\tilde{c} \cdot c).$$

Therefore, using Lemma 11 with $\tilde{c} := c'$ in the Step 3 we obtain

$$|X_{\text{forest}}| = O(c' \cdot c) \cdot 2k' = O((4c^3 + c) \cdot c \cdot (c^2 k \log^{3/2} n)) = O(c^6 k \log^{3/2} n).$$

Similarly, from Lemma 13, we have

$$|X_{\text{tripod}}| = O(c' \cdot c)O(k' \log n) = O((4c^3 + c) \cdot c) \cdot O(c^2 k \log^{3/2} n) \log n = O(c^6 k \log^{5/2} n),$$

which implies that

$$\kappa = O(ck \log^{3/2} n) + O(c^6 k \log^{3/2} n) + O(c^6 k \log^{5/2} n) = O(c^6 k \log^{5/2} n).$$

To find σ , we show that $G \setminus K$ admits a $O(c^4)$ -embedding ι into F'' with $\iota(v) = v$ for all $v \in G \setminus K$. By Lemma 18 F' is a c' -minor of $G' \setminus (X_{\text{forest}} \cup X_{\text{tripod}}) = G \setminus K$ with respect to the partition $\mathcal{P}' := \mathcal{P} \setminus (\cup_{v \in Y_{\text{forest}} \cup Y_{\text{tripod}}} \mathcal{P}(v))$. Consider arbitrary $x_1, x_2 \in V(G \setminus K)$ and let $v_1, v_2 \in V(F')$ be such that $x_1 \in \mathcal{P}'(v_1)$, $x_2 \in \mathcal{P}'(v_2)$. Let Q be the unique v_1 - v_2 path in F' . We use Q to construct a v_1 - v_2 path P in $G \setminus K$, with

$$\text{length}(Q) \leq \text{length}(P) \leq 2c'c \cdot \text{length}(Q).$$

Since F' is a c' -minor of $G \setminus K$, for any $\{w_1, w_2\} \in E(Q)$ there is $\{z_1, z_2\} \in E(G \setminus K)$ with $z_i \in \mathcal{P}'(w_i)$ for $i \in [2]$. Moreover, for any $w \in V(Q)$ the corresponding $\mathcal{P}'(w)$ is a connected subgraph such that $|V(\mathcal{P}'(w_i))| \leq 2c'\Delta(G \setminus K) + 1 = 2c'c + 1$. Thus, Q induces a walk $W \subseteq G \setminus K$ with $|V(W)| \leq 2c'c \cdot \text{length}(Q)$ and $v_1, v_2 \in W$. It follows that there is a v_1 - v_2 path P in W , such that

$$\text{length}(P) \leq 2c'c \cdot \text{length}(Q).$$

Note that since Q is the v_1 - v_2 shortest path in F' , we obtain

$$\text{length}(P) \leq 2c'c \cdot d_{F'}(v_1, v_2) = 2c'c \cdot d_{F''}(v_1, v_2),$$

where the last equality follows from the construction of F'' .

We claim that ι has contraction $O(c^4)$. By construction of F'' we have that $d_{F''}(x_i, v_i) = 1$ thus

$$d_{G \setminus K}(x_i, v_i) \leq c' \leq c' d_{F''}(x_i, v_i).$$

Therefore, we have that

$$\begin{aligned} d_{G \setminus K}(x_1, x_2) &\leq d_{G \setminus K}(x_1, v_1) + d_{G \setminus K}(v_1, v_2) + d_{G \setminus K}(v_2, x_2) \\ &\leq c' d_{F''}(x_1, v_1) + 2c'c \cdot \text{length}(Q) + c' d_{F''}(v_2, x_2) \\ &\leq 2c'c \cdot d_{F''}(x_1, v_1) + 2c'c \cdot d_{F''}(v_1, v_2) + 2c'c \cdot d_{F''}(v_2, x_2). \end{aligned}$$

Since F'' is a tree, it follows that

$$2c'c \cdot d_{F''}(x_1, v_1) + 2c'c \cdot d_{F''}(v_1, v_2) + 2c'c \cdot d_{F''}(v_2, x_2) = 2c'c \cdot d_{F''}(x_1, x_2).$$

Since $c' = O(c^3)$, it follows that the contraction of ι is at most $O(c^4)$. Now we prove that the expansion of ι is $O(1)$. We claim that $d_{F''}(x_1, x_2) \leq d_{G \setminus K}(x_1, x_2) + 2$. By the construction of F'' we have

$$\begin{aligned} d_{F''}(x_1, x_2) &= d_{F''}(x_1, v_1) + d_{F''}(v_1, v_2) + d_{F''}(v_2, x_2) \\ &= d_{F''}(x_1, v_1) + d_{F'}(v_1, v_2) + d_{F''}(v_2, x_2) = d_{F'}(v_1, v_2) + 2. \end{aligned}$$

Since F' is a c' -minor of $G \setminus K$, by Lemma 19 we get

$$d_{F'}(v_1, v_2) + 2 \leq d_{G \setminus K}(v_1, v_2) + 2,$$

thus the expansion of ι is $O(1)$. Therefore, the distortion of ι is $O(c^4)$. Hence, we obtain that the map $\phi := f \circ \iota : G \setminus K \rightarrow \mathbb{R}^1$ has distortion $\sigma = O(c^9) \cdot O(c^4) = O(c^{13})$, which concludes the proof. \blacktriangleleft

References

- 1 Vineet Bafna, Piotr Berman, and Toshihiro Fujito. A 2-approximation algorithm for the undirected feedback vertex set problem. *SIAM Journal on Discrete Mathematics*, 12(3):289–297, 1999.
- 2 Karol Borsuk. Drei sätze über die n-dimensionale euklidische sphäre. *Fundamenta Mathematicae*, 20(1):177–190, 1933.
- 3 Mihai Bădoiu, Piotr Indyk, and Anastasios Sidiropoulos. Approximation algorithms for embedding general metrics into trees. In Nikhil Bansal, Kirk Pruhs, and Clifford Stein, editors, *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, pages 512–521. SIAM, 2007. URL: <http://dl.acm.org/citation.cfm?id=1283383.1283438>.
- 4 Mihai Bădoiu, Julia Chuzhoy, Piotr Indyk, and Anastasios Sidiropou. Embedding ultrametrics into low-dimensional spaces. In *Proceedings of the twenty-second annual symposium on Computational geometry*, pages 187–196, 2006.
- 5 Mihai Bădoiu, Julia Chuzhoy, Piotr Indyk, and Anastasios Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 225–233, 2005.
- 6 Mihai Bădoiu, Kedar Dhamdhere, Anupam Gupta, Yuri Rabinovich, Harald Räcke, Ramamoorthi Ravi, and Anastasios Sidiropoulos. Approximation algorithms for low-distortion embeddings into low-dimensional spaces. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 119–128. Society for Industrial and Applied Mathematics, 2005.
- 7 Timothy Carpenter, Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Anastasios Sidiropoulos. Algorithms for low-distortion embeddings into arbitrary 1-dimensional spaces. In Bettina Speckmann and Csaba D. Tóth, editors, *34th International Symposium on Computational Geometry, SoCG 2018, June 11-14, 2018, Budapest, Hungary*, volume 99 of *LIPICs*, pages 21:1–21:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.SoCG.2018.21.
- 8 Victor Chepoi, Feodor F Dragan, Ilan Newman, Yuri Rabinovich, and Yann Vaxes. Constant approximation algorithms for embedding graph metrics into trees and outerplanar graphs. *Discrete & Computational Geometry*, 47(1):187–214, 2012.
- 9 Vasek Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of operations research*, 4(3):233–235, 1979.
- 10 Mark de Berg, Krzysztof Onak, and Anastasios Sidiropoulos. Fat polygonal partitions with applications to visualization and embeddings. *arXiv preprint*, 2010. arXiv:1009.1866.
- 11 Uriel Feige, MohammadTaghi Hajiaghayi, and James R Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM Journal on Computing*, 38(2):629–657, 2008.
- 12 Michael R. Fellows, Fedor V. Fomin, Daniel Lokshtanov, Elena Losievskaja, Frances A. Rosamond, and Saket Saurabh. Distortion is fixed parameter tractable. In Susanne Albers, Alberto Marchetti-Spaccamela, Yossi Matias, Sotiris E. Nikolettseas, and Wolfgang Thomas, editors, *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I*, volume 5555 of *Lecture Notes in Computer Science*, pages 463–474. Springer, 2009. doi:10.1007/978-3-642-02927-1_39.
- 13 Ian Goodfellow, Patrick McDaniel, and Nicolas Papernot. Making machine learning robust against adversarial inputs. *Communications of the ACM*, 61(7):56–66, 2018.
- 14 Piotr Indyk, Jiří Matoušek, and Anastasios Sidiropoulos. 8: low-distortion embeddings of finite metric spaces. In *Handbook of discrete and computational geometry*, pages 211–231. Chapman and Hall/CRC, 2017.
- 15 Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Slightly superexponential parameterized problems. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 760–776. SIAM, 2011. doi:10.1137/1.9781611973082.60.

- 16 Jiří Matoušek and Anastasios Sidiropoulos. Inapproximability for metric embeddings into \mathbb{R}^d . In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 405–413. IEEE Computer Society, 2008. doi:10.1109/FOCS.2008.21.
- 17 Amir Nayyeri and Benjamin Raichel. Reality distortion: Exact and approximate algorithms for embedding into the line. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 729–747. IEEE, 2015.
- 18 Amir Nayyeri and Benjamin Raichel. A treehouse with custom windows: Minimum distortion embeddings into bounded treewidth graphs. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 724–736. SIAM, 2017.
- 19 Anastasios Sidiropoulos, Dingkang Wang, and Yusu Wang. Metric embeddings with outliers. In Philip N. Klein, editor, *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 670–689. SIAM, 2017. doi:10.1137/1.9781611974782.43.

A Density Reduction

A.1 The algorithm for density reduction

Let us describe the algorithm for reducing the density of a graph. The algorithm takes as input a graph G and some $k \geq 0$, $c \geq 1$, and outputs some $Y \subseteq V(G)$, such that $\Delta(G \setminus Y) \leq c$. This is summarized in Algorithm 1.

■ **Algorithm 1** SPARSIFY.

```

1: procedure SPARSIFY( $G, c$ )
2:   if  $\Delta(G) \leq c$  then
3:     return  $\emptyset$ 
4:   else
5:     Let  $X$  be a  $3/4$ -balanced vertex separator of  $G$  computed by Theorem 5.
6:     Let  $G_1, \dots, G_t$  be the connected components of  $G \setminus X$ .
7:     return  $X \cup \left( \bigcup_{i=1}^t \text{SPARSIFY}(G_i, c) \right)$ 

```

A.2 Analysis of the algorithm for density reduction

We now analyze the algorithm described above. We first recall the following result from [6].

► **Lemma 20** (Bădoiu et al. [6]). *If G admits a c -embedding into the line then $\Delta(G) \leq c$.*

The following establishes the existence of small balanced separators.

► **Lemma 21.** *Let G be a graph such that G admits a (k, c) -embedding into the line. Let $Z \subseteq V(G)$ with $|Z| = k$ be such that $G \setminus Z$ is c -embeddable into the line. Then any $H \subseteq G$ contains a $2/3$ -balanced vertex separator of size at most $c + |Z \cap V(H)|$.*

Proof. Let $f : G \setminus Z \rightarrow \mathbb{R}$ be an embedding with distortion c . Let $V(H) = \{v_1, \dots, v_h\}$, and assume w.l.o.g. that $f(v_1) < f(v_2) < \dots < f(v_h)$. Let $X = \{v_{\lfloor h/3 \rfloor + 1}, v_{\max\{h, \lfloor h/3 \rfloor + c + 1\}}\}$. By lemma 16 we get that X is a balanced separator of $H \setminus Z = H \setminus (Z \cap V(H))$. Therefore, $W := X \cup (V(H) \cap Z)$ is a balanced separator for H , with $|W| = |X| + |Z \cap V(H)| \leq c + |Z \cap V(H)|$, as required. ◀

We are now ready to prove the main result of this Section.

36:18 Computing Bi-Lipschitz Outlier Embeddings into the Line

Proof of Lemma 10. It is immediate that the output, Y , of the procedure SPARSIFY is such that $\Delta(G \setminus Y) \leq c$. Also, if $\Delta(G) \leq c$, the algorithm outputs $Y = \emptyset$.

It thus remains to bound $|Y|$. Fix some $K \subseteq V(G)$, with $|K| = k$, and some c -embedding f of $G \setminus K$ into the line. Consider some recursive call of procedure SPARSIFY(H, c), for some $H \subseteq G$. If $H \cap K = \emptyset$, then $H \subseteq G \setminus K$, and thus $\Delta(H) \leq \Delta(G \setminus K) \leq c$, where the last inequality follows by lemma 20. Therefore, procedure SPARSIFY computes a balanced separator, X_H , only if H intersects K . By lemma 21 and Theorem 5 it follows that

$$|X_H| \leq O\left(\sqrt{\log n} \cdot (c + |K \cap V(H)|)\right) \leq O\left(|K \cap V(H)| \cdot c \cdot \sqrt{\log n}\right).$$

We charge the vertices in X_H to the vertices in $K \cap H$; thus every vertex in $K \cap H$ receives at most $O(\sqrt{\log n})$ units of charge. Since any two subgraphs on the same level of the recursion are disjoint, it follows that each vertex in K receives at most $O(c\sqrt{\log n})$ units of charge per level of the recursion. Since each separator is $3/4$ -balanced, it follows that the depth of the recursion is at most $\log_{4/3} n$. Thus, every vertex in K receives at most $\log_{4/3} n \cdot O(c\sqrt{\log n}) = \beta \cdot c \log_{4/3}^{3/2} n$ units of charge throughout the execution of the procedure SPARSIFY. The constant β comes from the bound on the size of the vertex separator computed by Theorem 5. Hence, if $Y > \beta \cdot kc \log_{4/3}^{3/2} n$, then we have certified that G does not admit a (k, c) -embedding into the line, which concludes the proof. ◀

B Eliminating large metrical cycles

B.1 The algorithm

The input consists of a graph G , some $c \geq 1$, and $k \geq 0$. The algorithm proceeds in steps, that are formally described below.

Algorithm for eliminating large metrical cycles:

Step 1. Compute a c -net N of G .

Step 2. Compute a Voronoi partition \mathcal{P} of G centered at N , and the corresponding c -minor H of G .

Step 3. Using the algorithm from Theorem 4 compute a 2-approximate solution S to the Minimum Feedback Vertex Set problem on H . If $|S| > 2k$, then decide that G does not admit a (k, c) -embedding into the line.

B.2 Analysis

First, we prove the following statement about embeddability into a subgraph of a c -minor.

► **Lemma 22.** *Let G be a graph, $R > 0$, let N be a R -net in G , let \mathcal{P} be a R -partition centered at N , and let H be the R -minor of G induced by \mathcal{P} . Let $X \subset N$, and let*

$$Y = \bigcup_{x \in X} \mathcal{P}(x).$$

Then the metric space $(N \setminus X, d_{G \setminus Y})$ admits a $(2R + 1)$ -embedding into $H \setminus X$. Moreover, this embedding can be computed in polynomial time.

Proof. Let $u, v \in N \setminus X$. Let Q be a u - v shortest path in $G \setminus Y$. When traversing Q starting from u let C_1, \dots, C_ℓ be the sequence of clusters of \mathcal{P} visited. For each $i \in [\ell]$ let q_i be the center of C_i ; that is, $C_i = \mathcal{P}(q_i)$. Since for all $i \in [\ell - 1]$ there is an edge in $G \setminus Y$ between some vertex in C_i and some vertex in C_{i+1} , it follows that there also exists an edge in $H \setminus X$ between q_i and q_{i+1} . Therefore $Q' = q_1, \dots, q_\ell$ is a path in $H \setminus X$. We thus obtain

$$d_{H \setminus X}(u, v) \leq \text{length}(Q') \leq \text{length}(Q) = d_{G \setminus Y}(u, v). \quad (4)$$

Let $W = w_1, \dots, w_t$ be a u - v shortest path in $H \setminus X$. Since each cluster in \mathcal{P} has radius at most R , it follows that for all $i \in [t-1]$ there exists a w_i - w_{i+1} path in $G \setminus Y$ of length at most $2R+1$. Concatenating all these paths we obtain a u - v path W' in $G \setminus Y$ of length at most $(t-1) \cdot (2R+1)$. Thus

$$d_{G \setminus Y}(u, v) \leq \text{length}(W') \leq (2R+1)(t-1) = (2R+1)d_{H \setminus X}(u, v). \quad (5)$$

Combining (4) and (5) the assertion follows. \blacktriangleleft

We recall the Borsuk-Ulam Theorem [2].

► Theorem 23 (Borsuk-Ulam Theorem [2]). *Let $d \geq 1$, and let \mathbb{S}^d denote the d -dimensional sphere. Let $f : \mathbb{S}^d \rightarrow \mathbb{R}^d$ be a continuous map. Then there exists $x \in \mathbb{S}^d$, such that $f(x) = f(-x)$.*

The following is a simple consequence of Theorem 23. A similar argument is used in [6].

► Lemma 24. *Let C be a cycle and let $f : V(C) \rightarrow \mathbb{R}$ be an injective map. Then there exist $u, v, w \in V(C)$, such that $\{u, v\} \in E(C)$, and $f(u) < f(w) < f(v)$.*

Proof. Suppose that C is the n -cycle for some $n \in \mathbb{N}$. We identify the vertices in C with distinct points in \mathbb{S}^1 , so that the points appear in the same order as in C along a clockwise traversal of \mathbb{S}^1 . For each $\{x, y\} \in E(C)$ there exists an arc $A_{x,y}$ in \mathbb{S}^1 that does not contain any other vertex in C ; we extend f to $A_{x,z}$ affinely. After repeating for all edges in C , we obtain a continuous map $f : \mathbb{S}^1 \rightarrow \mathbb{R}^1$. By Theorem 23 we get that there exists $x \in \mathbb{S}^1$ with $f(x) = f(-x)$. This means that there exist two edges in C whose images in f span overlapping intervals in \mathbb{R}^1 . Since f is injective on $V(C)$ this implies that one endpoint is contained inside the interval of the other edge, which concludes the proof. \blacktriangleleft

We next establish the existence of a small feedback vertex set in the minor computed by the algorithm.

► Lemma 25. *Let G be a graph, $c \geq 1$, $k \geq 0$, such that G admits a (k, c) -embedding into the line. Let H be a R -minor of G , for some $R \geq c$. Then there exists a feedback vertex set X in H with $|X| \leq k$.*

Proof. Let \mathcal{P} be the R -partition of G such that H is the R -minor of G induced by \mathcal{P} . Since G admits a (k, c) -embedding into the line, it follows that there exists some $Y \subseteq V(G)$, with $|Y| \leq k$, such that $G \setminus Y$ admits a c -embedding f into the line.

Let X be the set of all $v \in V(H)$, such that Y intersects the cluster in \mathcal{P} centered at v ; that is $X = \{v \in V(H) : \mathcal{P}(v) \cap Y \neq \emptyset\}$. Since \mathcal{P} is a partition, it is immediate that $|X| \leq |Y| \leq k$. It therefore remains to show that $H \setminus X$ is acyclic. Suppose, for the sake of contradiction, that $H \setminus X$ is not acyclic. Let C be a cycle in $H \setminus X$. By Lemma 24 there exist $u, v, w \in V(C)$, such that $\{u, v\} \in E(C)$, and $f(u) < f(w) < f(v)$.

Since $\{u, v\} \in E(C)$, and $C \subseteq H$, it follows that $\{u, v\} \in E(H)$. Since H is R -minor, it follows that there exists a path Q between u and v , with $Q \subseteq \mathcal{P}(u) \cup \mathcal{P}(v)$. When traversing Q starting from u let u' be the last vertex visited with $f(u') < f(w)$; let also v' be the vertex visited immediately after u' . We have $f(u') < f(w) < f(v')$.

36:20 Computing Bi-Lipschitz Outlier Embeddings into the Line

Since H is a R -minor and $u' \notin \mathcal{P}(w)$, it follows that $d_G(w, u') \geq d_G(u, u')$. By the definition of a R -partition we have that $d_G(u, w) > c$, and therefore $d_G(u', w) > R/2$. Similarly, we obtain $d_G(v', w) > R/2$. Since f is non-contracting, we obtain $|f(u') - f(v')| = |f(u') - f(w)| + |f(w) - f(v')| \geq d(u, w) + d(w, v) > R/2 + R/2 = R \geq c$, which contradicts the fact that f has expansion at most c , and concludes the proof. \blacktriangleleft

We are now ready to prove the main result of this Section.

Proof of Lemma 11. By Lemma 25, either G does not admit a (k, c) -embedding into the line, or there exists $X \subseteq V(H)$, with $|X| \leq k$, such that $H \setminus X$ is acyclic. Using the algorithm from Theorem 4 we compute in Step 3 a 2-approximation $S \subseteq V(H)$ to the Minimum Feedback Vertex Set in H . Therefore, if $|S| > 2k$, then we can terminate with outcome (1), and otherwise terminate with outcome (2), which completes the proof. \blacktriangleleft

C Eliminating large metrical tripods

In this Section we present and analyze the procedure for eliminating large metrical tripods. We begin by showing that large tripods are an obstruction to embeddability into the line. This is summarized in Lemma 12.

Proof of Lemma 12. Let f be a non-contractive embedding of J into the line. Let v be the common endpoint of P_1, P_2, P_3 . For each $i \in [3]$ let v_i be the other endpoints of P_i . We may assume w.l.o.g. (by change of indices) that $f(v_1) < f(v_2) < f(v_3)$. Let Q be the unique v_1 - v_3 path in J . It follows that there exists $\{u, w\} \in E(Q)$, such that $f(u) < f(v_2) < f(w)$. This implies that $|f(u) - f(w)| = |f(u) - f(v_2)| + |f(v_2) - f(w)| \geq d_G(u, v_2) + d_G(v_2, w) \geq 2R = 2Rd_J(u, w)$. Therefore the distortion of f is at least $2R$, which concludes the proof. \blacktriangleleft

The above easily implies the following results, which asserts the existence of a small set of vertices whose removal eliminates all large tripods.

► Lemma 26. *Let F be a forest that admits a (k, c) -embedding into the line. Then there exists some $X \subseteq V(F)$, with $|X| \leq k$, such that $F \setminus X$ does not contain any $(c/2 + 1)$ -tripod as a subgraph.*

Proof. Since F admits a (k, c) -embedding into the line, it follows that there exists some $X \subseteq V(F)$, with $|X| \leq k$, such that $F \setminus X$ admits a c -embedding into the line. It suffices to show that $F \setminus X$ does not contain any $(c/2 + 1)$ -tripods. Suppose, for the sake of contradiction, that $F \setminus X$ contains some $(c/2 + 1)$ -tripod J . Since $(V(J), d_J)$ is a submetric of $(V(F) \setminus X, d_{F \setminus X})$, it follows that J admits a c -embedding into the line, which contradicts Lemma 12, and concludes the proof. \blacktriangleleft

Now are now ready to prove the main result of this Section.

Proof of Lemma 13. Any tripod $T \subseteq F$ can be uniquely specified by selecting its root and its three leaves. Therefore, there are at most $O(|V(F)|^4)$ distinct tripods in F . Moreover, the set of all tripods, \mathcal{T} , can be enumerated in polynomial time. We form an instance of the Minimum Set Cover problem with universe $U = \mathcal{T}$. We also let

$$\mathcal{C} = \bigcup_{v \in V(F)} \{C_v\},$$

where $C_v = \{T \in \mathcal{T} : v \in V(T)\}$. It is immediate that for any $Y \subseteq V(F)$, $F \setminus Y$ contains no R -tripods iff $\bigcup_{v \in Y} C_v = U$. Therefore, computing a minimum-cardinality subset of vertices of F whose deletion removes all R -tripods, is equivalent to solving the Minimum Set Cover instance on (U, \mathcal{C}) . The result now follows from Theorem 6. \blacktriangleleft

D Embedding Trees Without Large Tripods into the Line

This Section is devoted to proving Lemma 14, which asserts that any tree with no large tripods admits a low-distortion embedding into the line.

Proof of Lemma 14. Since T is a tree, we can compute in polynomial time a longest path Q in T . Let $Q = v_1, \dots, v_t$. Let \mathcal{P} be a Voronoi partition centered at $V(Q)$. Since T does not contain any R -tripod as a subgraph, it follows that for all $u \in V(T)$, there exists some $v \in V(Q)$, with $d_T(u, v) < R$. Therefore, for each $v_i \in V(Q)$, we have

$$|\mathcal{P}(v_i)| \leq |\text{Ball}_T(v_i, R-1)| \leq \Delta(T) \cdot 2(R-1) + 1 \leq \Delta(T) \cdot 2R - 1.$$

By the definition of a graphical Voronoi partition we have that for all $i \in [t]$, the vertex-induced subgraph $T_i := T[\mathcal{P}(v_i)]$ is connected, and thus T_i is a subtree of T . Let W_i be a closed walk in T_i that visits all vertices in T_i , obtained by duplicating every edge (or, equivalently, the walk obtained by any traversal of T_i). Since every edge in T_i is traversed twice, we have $\text{length}(W_i) = 2(|V(T_i)| - 1)$. Let $W_i = w_{i,1}, \dots, w_{i,t_i}$.

We define the embedding $f_i : V(T_i) \rightarrow \mathbb{R}$ as follows. For each $v \in V(T_i)$, we define $f_i(v) = \min\{j \in [t_i] : v = w_{i,j}\}$.

We combine the mappings f_1, \dots, f_t into a mapping $f : V(G) \rightarrow \mathbb{R}$. Informally, this is done by translating each f_i so that for all $i \in [t-1]$, the image of f_i appears to the left of the image of f_{i+1} , and there is a gap of length $2R$ between these two images.

Formally, for each $u \in \mathcal{P}_{v_i}$, we set $f(u) = L_i + f_i(u)$, where

$$L_i = \begin{cases} 0 & \text{if } i = 0 \\ L_{i-1} + \max_{z \in \mathcal{P}(v_{i-1})} \{f_{i-1}(z)\} + 2R & \text{otherwise} \end{cases}$$

This completed the definition of the embedding f .

It remains to bound the distortion of f . For vertices that lie in the same cluster in \mathcal{P} , the map is non-contractive since the distance in the embedding is at least the distance in some walk W_i , which is at least the distance in T . Moreover, the expansion is upper bounded by the length of the walk, which is at most $\Delta(T) \cdot (2R - 1)$.

Next, let us consider $p, q \in V(T)$ that fall in different clusters in \mathcal{P} . Suppose that $p \in \mathcal{P}(v_i)$, and $q \in \mathcal{P}(v_j)$, for some $i, j \in [t]$, with $i < j$. We have

$$\begin{aligned} |f(p) - f(q)| &\leq 2R(j-1) + \sum_{r=i}^j \text{length}(W_r) \leq (j-1)2R + (j-i+1)\Delta(T) \cdot (2R-1) \\ &\leq (j-i) \cdot O(\Delta(T) \cdot R) = d_T(v_i, v_j) \cdot O(\Delta(T) \cdot R) \leq d_T(p, q) \cdot O(\Delta(T) \cdot R). \end{aligned}$$

Moreover $|f(p) - f(q)| \geq 2R(j-i) + 1 \geq 2R + (j-i) \geq d_T(p, v_i) + d_T(v_i, v_j) + d_T(v_j, q) = d_T(p, q)$.

Therefore, in all cases, f is non-contractive and has expansion at most $O(\Delta(T) \cdot R)$. \blacktriangleleft