

Lower Bounds for Embedding into Distributions over Excluded Minor Graph Families

Douglas E. Carroll ^{*1} and Ashish Goel^{**2}

¹ UCLA

² Stanford University

Abstract. It was shown recently by Fakcharoenphol *et al.* [7] that arbitrary finite metrics can be embedded into distributions over tree metrics with distortion $O(\log n)$. It is also known that this bound is tight since there are expander graphs which cannot be embedded into distributions over trees with better than $\Omega(\log n)$ distortion.

We show that this same lower bound holds for embeddings into distributions over any minor excluded family. Given a family of graphs F which excludes minor M where $|M| = k$, we explicitly construct a family of graphs with treewidth- $(k + 1)$ which cannot be embedded into a distribution over F with better than $\Omega(\log n)$ distortion. Thus, while these minor excluded families of graphs are more expressive than trees, they do not provide asymptotically better approximations in general. An important corollary of this is that graphs of treewidth- k cannot be embedded into distributions over graphs of treewidth- $(k - 3)$ with distortion less than $\Omega(\log n)$.

We also extend a result of Alon *et al.* [1] by showing that for any k , planar graphs cannot be embedded into distributions over treewidth- k graphs with better than $\Omega(\log n)$ distortion.

1 Introduction

Many difficult problems can be approximated well when restricted to certain classes of metrics [13, 3, 4]. Therefore, low distortion embeddings into these restricted classes of metrics become very desirable. We will refer to the original metric which we would like to embed as the *source metric* and the metric into which we would like to embed as the *target metric*.

Tree metrics form one such class of desirable target metrics in the sense that many difficult problems become tractable when restricted to tree metrics. However, they are not sufficiently expressive; i.e. it has been shown that there are classes of metrics which cannot be approximated well by trees. In particular,

* Department of Computer Science, University of California, Los Angeles. Email: dcarroll@cs.ucla.edu.

** Departments of Management Science and Engineering and (by courtesy) Computer Science, Stanford University. Research supported in part by NSF CAREER award 0133968, and by a Terman Engineering Fellowship. Email: ashishg@stanford.edu.

Rabinovich and Raz [15] have proved that graph metrics cannot be embedded into trees with distortion better than $\text{girth}/3 - 1$.

Therefore, subsequent approaches have proposed the use of more expressive classes of metrics. Alon *et al.* [1] showed that any n -point metric can be embedded with $2^{O(\sqrt{\log n \log \log n})}$ distortion into distributions over spanning trees. In so doing they demonstrated that such probabilistic metrics are more expressive than tree metrics. Bartal [2] formally defined probabilistic embeddings and proposed using distributions over arbitrary dominating tree metrics. He showed that any finite metric could be embedded into distributions over trees with $O(\log^2 n)$ distortion. He subsequently improved this bound to $O(\log n \log \log n)$ [3].

This path culminated recently in the result of Fakcharoenphol *et al.* [7] in which they improved this bound to $O(\log n)$ distortion. This upper bound is known to be tight since there exist graphs which cannot be embedded into such distributions with better than $\Omega(\log n)$ distortion. This lower bound follows naturally from the fact that any distribution over trees can be embedded into ℓ_1 with constant distortion, and the existence of expander graphs which cannot be embedded into ℓ_1 with distortion better than $\Omega(\log n)$. Surprisingly, Gupta *et al.* [8] showed that this same bound is in fact achieved by source graph metrics of treewidth-2. Their result is also implicit in the work of Imase and Waxman [10] where they establish a lower bound on the competitive ratio of online Steiner trees.

It is plausible to hypothesize that there are more general and expressive classes of target metrics. We explore using distributions over minor closed families of graphs and show that asymptotically, they are no stronger than distributions over trees. We show a $\Omega(\log n)$ lower bound on the distortion even when the source metrics are graphs of low treewidth. More precisely, for any minor M , where $|M| = k$, we exhibit a construction for an infinite class of finite metrics of treewidth- $(k + 1)$ for which any embedding into distributions over families of graphs excluding M achieves $\Omega(\log n)$ distortion. A corollary of this fact is that treewidth- k graphs cannot be embedded into distributions over treewidth- $(k - 3)$ graphs with distortion less than $\Omega(\log n)$.

A weaker result can be inferred directly from Rao [16] who proved that any minor closed family can be embedded into ℓ_1 with distortion $O(\sqrt{\log n})$. Consequently, the expanders exhibited by Linial *et al.* [13] cannot be embedded into distributions over minor closed families with distortion less than $\Omega(\sqrt{\log n})$.

One can derive a lower bound of $\Omega(\log n)$ by combining the method of Klein *et al.* [11] for decomposing minor excluded graphs with Bartal's [2] proof of the lower bound for probabilistic embeddings into trees. This bound also follows from the recent paper of Rabinovich on the average distortion of embeddings into ℓ_1 . However, we show this same bound holds for embeddings of simple, low-treewidth source metrics.

We continue by exploring the difficulty of embedding planar graph metrics into distributions over other minor excluded graph families. Alon *et al.* showed that any embedding of a 2-dimensional grid into a distribution over spanning subtrees must have distortion $\Omega(\log n)$. We show that for any fixed k , the same

lower bound holds for embeddings of 2-dimensional grids into distributions over dominating treewidth- k graph metrics.

Note that our $\Omega(\log n)$ lower bounds hide polynomial factors of k .

1.1 Techniques

We employ Yao’s MiniMax principle to prove both lower bounds – it suffices to show that for some distribution over the edges of the source graph, any embedding of the source graph into a dominating target graph has large expected distortion. In both cases the expected distortion is shown to be logarithmic in the number of vertices in the graph.

For the main result, we show that given any minor M , one can recursively build a family of source graphs which guarantee a large expected distortion when embedded into a graph which excludes M as a minor. This structure is detailed in section 3. While we use some elements from the construction of Gupta *et al.* [8], the main building block in our recursive construction is a stretched clique rather than the simpler treewidth-2 graphs used by Gupta *et al.*, and hence, our proof is more involved. As a step in our proof, we show that if graph G does not contain H as a minor, then subdivisions of H can not be embedded into G with a small distortion.

To show the lower bound for embedding planar graphs into bounded tree-width graphs, we use the notion of nice tree decompositions [5, 12] of bounded tree-width graphs. We show that for a range of set sizes, the subsets of a 2-dimensional grid have much larger separators than do their embeddings in the nice tree decompositions; this suffices to prove an $\Omega(\log n)$ lower bound with a little more work. This argument is presented in section 4.

2 Definitions and Preliminaries

Given two metric spaces (G, ν) and (H, μ) and an embedding $\Phi : G \rightarrow H$, we say that the *distortion* of the embedding is $\|\Phi\| \cdot \|\Phi^{-1}\|$ where

$$\|\Phi\| = \max_{x, y \in G} \frac{\mu(\Phi(x), \Phi(y))}{\nu(x, y)},$$

$$\|\Phi^{-1}\| = \max_{x, y \in G} \frac{\nu(x, y)}{\mu(\Phi(x), \Phi(y))}$$

and (G, ν) α -*approximates* (H, μ) if the *distortion* is no more than α . We say that μ *dominates* ν if $\mu(\Phi(x), \Phi(y)) \geq \nu(x, y) \forall x, y$.

Definition 1. Given a graph $G = (V_G, E_G)$, a tree $T = (V_T, E_T)$ and a collection $\{X_i | i \in V_T\}$ of subsets of V_G , then $(\{X_i\}, T)$ is said to be a tree decomposition of G if

1. $\bigcup_{i \in V_T} X_i = V_G$,

2. for each edge $e \in E_G$, there exists $i \in V_T$ such that the endpoints of e are in X_i , and
3. for all $i, j, k \in V_T$: if j lies on the path from i to k in T , then $X_i \cap X_k \subseteq X_j$.

The width of a tree decomposition is $\max_{i \in V_T} |X_i| - 1$. The treewidth of a graph G is the minimum width over all tree decompositions of G .

Definition 2. A tree decomposition $(\{X_i\}, T = (V_T, E_T))$ of a graph $G = (V_G, E_G)$ is said to be a nice decomposition if:

1. T is a rooted binary tree
2. if node $i \in V_T$ has two children j_1, j_2 , then $X_{j_1} = X_{j_2} = X_i$
3. if node $i \in V_T$ has one child j , then either $X_i \subset X_j$ and $|X_i| = |X_j| - 1$ or $X_j \subset X_i$ and $|X_j| = |X_i| - 1$
4. if node $i \in V_T$ is a leaf, then $|X_i| = 1$

Proposition 1. [5, 12] If graph G has a tree decomposition of width k , then G has a nice tree decomposition of width k .

The following property follows directly from the definition of a tree decomposition:

Proposition 2. For all $i, j, k \in V_T$: if j lies on the path from i to k in T , and $x_1 \in X_i$ and $x_2 \in X_k$ then any path connecting x_1 and x_2 in G must contain a node in X_j .

For a more complete exposition of treewidth see [12, 5].

Definition 3. Given a graph H , a k -subdivision of H is defined as the graph resulting from the replacement of each edge by a path of length k .

We will use $G_{n,m}$ to denote the m -subdivision of K_n . We observe that $\text{treewidth}(G_{n,m}) = n - 1 = \text{treewidth}(K_n)$.

3 Distributions over Minor Excluded Families

3.1 Outline

Given a family F which excludes minor M where $|M| = k$, we recursively construct a family of graphs H_i which embed into any graph in F with average distortion $\Omega(\log n)$ where $n = |H_i|$. Then by Yao's MiniMax principle, any embedding into a distribution over F must have distortion $\Omega(\log n)$ as claimed.

3.2 Results

Before proceeding with our main result, we need the following technical lemma³.

Lemma 1. *Let G, H be graphs such that G does not contain H as a minor and let J be the k -subdivision of H . If f is an embedding of the metric of J into the metric of G , then f has distortion at least $k/6 - 1/2$.*

Proof. We will use the linear extension framework of Rabinovich and Raz [15]. While we borrow some of their techniques, we have not found a reduction that would use their lower bounds (on the distortion of an embedding of a graph into another with a smaller Euler characteristic) as a black-box to prove this technical lemma.

The linear extension \tilde{Q} of a graph Q is obtained by identifying with each edge of Q a line of length 1. All the points on all the lines belong to \tilde{Q} . The metric d_Q can be extended to the metric \tilde{d}_Q in the following natural fashion. If x and y lie on the same line (u, v) then $\tilde{d}_Q(x, y)$ is merely their distance on the line. If x lies on the line (u, v) and y lies on a different line (w, r) , then

$$\begin{aligned} \tilde{d}_Q(x, y) = \min\{ & \tilde{d}_Q(x, u) + d_Q(u, w) + \tilde{d}_Q(w, y), \\ & \tilde{d}_Q(x, v) + d_Q(v, w) + \tilde{d}_Q(w, y), \\ & \tilde{d}_Q(x, u) + d_Q(u, r) + \tilde{d}_Q(r, y), \\ & \tilde{d}_Q(x, v) + d_Q(v, r) + \tilde{d}_Q(r, y)\} \end{aligned}$$

We will refer to the original vertices as vertices of \tilde{Q} and to the original edges of Q as the lines of \tilde{Q} . We can now extend the embedding f into a continuous map $\tilde{f} : \tilde{J} \rightarrow \tilde{G}$ such that

1. \tilde{f} agrees with f on vertices of \tilde{J} , and
2. if $x \in (u, v)$ where (u, v) is a line of \tilde{J} then $\tilde{f}(x)$ lies on a shortest path from $\tilde{f}(u)$ to $\tilde{f}(v)$ in \tilde{G} such that $\tilde{d}_J(x, u)/\tilde{d}_J(x, v) = \tilde{d}_G(\tilde{f}(x), \tilde{f}(u))/\tilde{d}_G(\tilde{f}(x), \tilde{f}(v))$.

Since \tilde{f} is continuous, the entire line (u, v) in \tilde{J} must be mapped to a single shortest path from $\tilde{f}(u)$ to $\tilde{f}(v)$ in \tilde{G} . We will now assume that the distortion α of f is less than $k/6 - 1/2$ and derive a contradiction. But first, we need to state and prove the following useful claim:

Claim. If the points $\tilde{f}(x)$ and $\tilde{f}(y)$ lie on the same line in \tilde{G} , the points x, x' lie on the same line in \tilde{J} , and the points y, y' lie on the same line in \tilde{J} , then $\tilde{d}_J(x', y') \leq 2\alpha + 1$.

Proof. Suppose x, x' and y, y' lie on lines (p, q) and (r, s) in \tilde{J} , respectively. Use X to denote the quantity $\tilde{d}_G(\tilde{f}(x), \tilde{f}(y))$ and Y to denote the quantity $\tilde{d}_G(\tilde{f}(p), \tilde{f}(x)) + \tilde{d}_G(\tilde{f}(x), \tilde{f}(q)) + \tilde{d}_G(\tilde{f}(r), \tilde{f}(y)) + \tilde{d}_G(\tilde{f}(y), \tilde{f}(s))$. Since $\tilde{f}(x)$ and

³ This lemma appears to be folklore. But we have not been able to find a proof in the literature. Our proof is non-trivial and might be useful in other contexts, so we have provided the proof in this paper.

$\tilde{f}(y)$ lie on the shortest paths from $\tilde{f}(p)$ to $\tilde{f}(q)$ and from $\tilde{f}(r)$ to $\tilde{f}(s)$, respectively, we have $\tilde{d}_G(\tilde{f}(p), \tilde{f}(q)) + \tilde{d}_G(\tilde{f}(r), \tilde{f}(s)) = Y$. Also, by triangle inequality, we have $\tilde{d}_G(\tilde{f}(p), \tilde{f}(r)) + \tilde{d}_G(\tilde{f}(p), \tilde{f}(s)) + \tilde{d}_G(\tilde{f}(q), \tilde{f}(r)) + \tilde{d}_G(\tilde{f}(q), \tilde{f}(s)) \leq 4X + 2Y$. Clearly, $X \leq 1$ and $Y \geq 2$, so we have

$$\frac{\tilde{d}_G(\tilde{f}(p), \tilde{f}(r)) + \tilde{d}_G(\tilde{f}(p), \tilde{f}(s)) + \tilde{d}_G(\tilde{f}(q), \tilde{f}(r)) + \tilde{d}_G(\tilde{f}(q), \tilde{f}(s))}{\tilde{d}_G(\tilde{f}(p), \tilde{f}(q)) + \tilde{d}_G(\tilde{f}(r), \tilde{f}(s))} \leq 4.$$

Since the distortion is α , we must have

$$\frac{\tilde{d}_J(p, r) + \tilde{d}_J(p, s) + \tilde{d}_J(q, r) + \tilde{d}_J(q, s)}{\tilde{d}_J(p, q) + \tilde{d}_J(r, s)} \leq 4\alpha.$$

But $\tilde{d}_J(p, q) = \tilde{d}_J(r, s) = 1$. Also, $\tilde{d}_J(p, r) + \tilde{d}_J(p, s) + \tilde{d}_J(q, r) + \tilde{d}_J(q, s) \geq 4\tilde{d}_J(x', y') - 4$. Hence, we have $(4\tilde{d}_J(x', y') - 4)/2 \leq 4\alpha$, or $\tilde{d}_J(x', y') \leq 2\alpha + 1$.

We will now continue with the proof of Lemma 1. For any edge $(u, v) \in H$, consider the u - v path of length k in \tilde{J} . Consider the image of this path in \tilde{G} under \tilde{f} . The image need not be a simple path, but must contain a simple path from $\tilde{f}(u)$ to $\tilde{f}(v)$. We choose any such simple path arbitrarily, and call it the representative path for (u, v) , and denote it by $P(u, v)$. Start traversing this path from $\tilde{f}(u)$ to $\tilde{f}(v)$ and stop at the first vertex q which is the image of a point x on the u - v path in \tilde{J} such that $\tilde{d}_J(u, x) \geq k/2 - \alpha - 1/2$. Let (p, q) be the last edge traversed. We will call this the representative line of (u, v) and denote it by $L(u, v)$. Consider a point y on the u - v path in \tilde{J} that maps to p . The choice of p and q implies that $\tilde{d}_J(u, y) < k/2 - \alpha - 1/2$, and hence, we can apply claim 1 to conclude that $k/2 - \alpha - 1/2 \leq \tilde{d}_J(u, x) \leq k/2 + \alpha + 1/2$. Now, for any point z not on the u - v path in \tilde{J} , we have $\tilde{d}_J(x, z) \geq k/2 - \alpha - 1/2$. Since we assumed $\alpha < k/6 - 1/2$, we have $\tilde{d}_J(x, z) > 2\alpha + 1$ and hence, $\tilde{f}(z)$ can not lie on the line $L(u, v)$.

Thus the representative line $L(u, v)$ has two nice properties: it lies on a simple path from $\tilde{f}(u)$ to $\tilde{f}(v)$ and does not contain the image of any point not on the u - v path in \tilde{J} .

Now, we perform the following simple process in the graph G : For every edge (u, v) of H , contract all the edges on the representative path $P(u, v)$ in G except the representative edge $L(u, v)$. Since the representative line of an edge (u, v) of H does not intersect the representative path of any other edge of H , no representative edge gets contracted. The resulting graph is a minor of G , but also contains H as a minor, which is a contradiction. Hence $\alpha \geq k/6 - 1/2$.

Now we can prove our main result:

Theorem 1. *Let F be a minor closed family of graphs which excludes minor M where $|V_M| = k$. There exists an infinite family of graphs H_i with treewidth- $(k+1)$ such that any α -approximation of the metric of H_i by a distribution over dominating graph metrics in F has $\alpha = \Omega(\log |H_i|)$.*

Proof. Assume k is odd. We proceed by constructing an infinite sequence H_i of graphs of treewidth- k and show that the minimum distortion with which they can be embedded into a distribution over graph metrics in F grows without bound as i increases.

First we construct H_1 : Construct the graph H_1 by taking K_k and attaching each point v to source s with $(k-1)/2$ disjoint paths of length k and attaching v to sink t with $(k-1)/2$ paths of length k . Call s and t the terminals of H_1 .

We now show that H_1 has $m = (2k+1)\binom{k}{2}$ edges and contains $\binom{k}{2}$ edge-disjoint s, t paths of length $2k+1$. Label the s to v_i paths $sp_{i,1}$ to $sp_{i,(k-1)/2}$ and the v_i to t paths $tp_{i,1}$ to $tp_{i,(k-1)/2}$. Observe that the paths formed as follows are edge disjoint:

path $sp_{i,j}$ followed by edge $(v_i, v_{(i+2j) \bmod k})$ followed by path $tp_{(i+2j) \bmod k, j}$

Note that H_1 has treewidth $\leq k+1$.

The H_i are constructed recursively: For each i , construct the graph H_i by replacing every edge in H_1 with a copy of H_{i-1} . Therefore, H_i has $(2k+1)^i \binom{k}{2}^i$ edges and each pair of non-terminal vertices is connected by $\binom{k}{2}^{i-1}$ edge-disjoint paths of length $(2k+1)^{i-1}$. Also note that H_i has treewidth $\leq k+1$.

As in [8] we use Yao's MiniMax Principle to prove the lower bound. We show that there is a distribution d over the edges of H_i such that for any embedding of H_i into a graph which excludes minor M , an edge chosen randomly from distribution d has an expected distortion of $\Omega(i)$. Then by Yao's MiniMax principle, for any embedding of H_i into a distribution over graphs in F , there must be an edge with expected distortion $\Omega(i)$. We shall assume a uniform distribution over the edges.

Let U be a graph which excludes minor M and let $\Phi : H_i \rightarrow U$ be an embedding of H_i into U such that distances in U dominate their corresponding distances in H_i . For each edge $e \in H_i$ we will give e the color $j, 1 \leq j \leq i-1$ if Φ distorts e by at least $\frac{1}{6}(2k+1)^j - 1/2$. Note that an edge may have many colors.

Consider the copies of $G_{k,(2k+1)^{i-1}}$ in H_i . H_i contains a copy of K_k in which every edge has been replaced with a copy of H_{i-1} . Each copy of H_{i-1} has $\binom{k}{2}^{i-1}$ edge disjoint paths of length $(2k+1)^{i-1}$. Thus, H_i clearly contains at least $\binom{k}{2}^{i-1}$ edge disjoint copies of $G_{k,(2k+1)^{i-1}}$.

Clearly, each copy of $G_{k,(2k+1)^{i-1}}$ contains K_k and therefore M as a minor since $|M| = k$. Thus, by Lemma 1, at least one edge in each copy of $G_{k,(2k+1)^{i-1}}$ has distortion $\geq \frac{1}{6}(2k+1)^{i-1} - 1/2$. Since H_i comprises m copies of H_{i-1} , it contains $m \binom{k}{2}^{i-2}$ copies of $G_{k,(2k+1)^{i-2}}$. Therefore, there are at least $m \binom{k}{2}^{i-2}$ edges with color $i-2$. In general, there are at least $\geq m^{i-j-1} \cdot \binom{k}{2}^j$ edges with color j .

The distortion of an edge is $\geq \frac{1}{6}(2k+1)^j - 1/2$ where j is the largest of the edge's colors. For each edge $e \in E_{H_i}$ let C_e be the set of colors which apply to edge e . Clearly, $\forall e \in E_{H_i}$,

$$\sum_{j \in C_e} \left(\frac{1}{6} (2k+1)^j - 1/2 \right) \leq 2 \cdot \max_{j \in C_e} \left(\frac{1}{6} (2k+1)^j - 1/2 \right)$$

Thus,

$$\begin{aligned} \sum_{e \in E_{H_i}} \max_{j \in C_e} \left(\frac{1}{6} (2k+1)^j - 1/2 \right) &\geq \frac{1}{2} \sum_{e \in E_{H_i}} \sum_{j \in C_e} \frac{1}{7} (2k+1)^j \\ &> \frac{1}{14} \sum_{j=1}^{i-1} |\{e | j \in C_e\}| \cdot (2k+1)^j \\ &\geq \frac{1}{14} \sum_{j=1}^{i-1} m^{i-j-1} \cdot \binom{k}{2}^j \cdot (2k+1)^j \\ &= \frac{1}{14} \sum_{j=1}^{i-1} m^{i-1} \\ &= \frac{1}{14} (i-1) \cdot m^{i-1} \\ &\geq \frac{1}{28m} i \cdot m^i \\ &\geq \frac{1}{28k^3} i \cdot m^i \end{aligned}$$

Then, since the H_i has m^i edges, the average distortion is: $\frac{1}{28k^3} i = \Omega(\log |H_i|)$.

4 Planar Graphs

4.1 Outline

First we show that given a 2-dimensional grid, there is a distribution over the edges such that any embedding into a treewidth- k graph metric has high expected distortion. The proof builds on the work of Alon *et al.* [1]. By Yao's Minimax principle, this is enough to show that the 2-dimensional grid can not be embedded into a distribution over such graphs with distortion less than $\Omega(\log n)$.

4.2 Results

In this section we will use $GRID_n$ to denote the planar graph consisting of the $n \times n$ grid.

Lemma 2. (From [1]) *If A is a set of β^2 vertices in $GRID_n$, where $\beta^2 \leq \frac{n^2}{2}$, then there are at least β rows that A intersects but does not fill or at least β columns that A intersects but does not fill.*

Lemma 3. (Modified from [1]) *If A is a set of β^2 vertices in $GRID_n$, where $\beta^2 \leq \frac{n^2}{2}$, and B is a set of at most $\beta/4$ vertices in A , then there are at least $\beta/2$ vertices in A that have neighbors outside A and have distance at least $\frac{\beta}{4|B|}$ from each vertex of B .*

Proof. By Lemma 2, there is a set C of at least β vertices in A which are in distinct rows or distinct columns and have neighbors outside of A . Since they are in distinct rows, a vertex of B can be at distance $< \frac{\beta}{4|B|}$ of at most $\frac{\beta}{2|B|}$ vertices in C . Thus, there are at least $\frac{\beta}{2}$ vertices at distance at least $\frac{\beta}{4|B|}$ from each vertex in B .

Lemma 4. *Let H be a graph of treewidth k , $\Phi : GRID_n \rightarrow H$ be an embedding of $GRID_n$ into H , and $\beta \leq n/4$. Then there are at least $\frac{n^2}{24\beta}$ edges (u, v) such that $d_H(u, v) > \frac{\beta}{16(k+1)}$*

Proof. Since H has treewidth- k , it must have a tree decomposition of width k . Moreover, it must have a nice tree decomposition (X_i, T) of width k by Proposition 1.

Given a subtree S of T , define the $H\text{SIZE}(S) = |\bigcup_{i \in V_S} X_i|$. Since (X_i, T) is a nice decomposition, we know that T has a maximum degree of 3. We also know that for any two adjacent vertices $i, j \in T$, $|X_i - X_j| \leq 1$. Thus, every subtree S of T has an edge e such that removing e creates 2 subtrees each with $H\text{SIZE}$ at least $1/3 \cdot H\text{SIZE}(S)$. Note that if S_1 and S_2 are the two subtrees of T , $\bigcup_{i \in V_{S_1}} X_i$ and $\bigcup_{i \in V_{S_2}} X_i$ are not disjoint.

Start with T and successively delete edges from the remaining component with the largest $H\text{SIZE}$ such that the $H\text{SIZE}$ s of the resulting subtrees are as evenly divided as possible. Do this until $\lceil \frac{n^2}{3\beta^2} \rceil - 1$ edges have been deleted and there are $\lceil \frac{n^2}{3\beta^2} \rceil$ pieces. The smallest piece will always be at least $1/3$ the $H\text{SIZE}$ of the previous largest piece. Therefore, on the average these pieces have $H\text{SIZE} = 3\beta^2$ and the smallest will have $H\text{SIZE} \geq \beta^2$.

Since each deleted edge of T is incident to 2 pieces, the average number of pieces incident with a piece is less than 2. Thus, at least half the pieces are incident with no more than 4 edges.

Each deleted edge in T represents a set of points which form a vertex cut of H of size $\leq k + 1$. Thus, there are $\frac{n^2}{6\beta^2}$ pieces of $H\text{SIZE} \geq \beta^2$ which are separated from the rest of H by a cut of size $\leq 4(k + 1)$. Let A be a piece of $H\text{SIZE} \geq \beta^2$ and let B be the subset of size $\leq 4(k + 1)$ separating A from the rest of H . Then by Lemma 3, A has at least $\beta/2$ vertices with neighbors outside of the piece whose distance from the vertices of B is at least $\frac{\beta}{16(k+1)}$. Thus, there are at least $\frac{n^2}{24\beta}$ edges which are each distorted by a factor of $\frac{\beta}{16(k+1)}$.

Theorem 2. *Any α -approximation of the metric of $GRID_n$ by a distribution over dominating treewidth k graphs has $\alpha = \Omega(\log n)$.*

Proof. Let H be an arbitrary graph of treewidth k whose metric dominates that of $GRID_n$. By Lemma 4, there are at least $\frac{n^2}{24\beta}$ edges which are distorted by $> \beta/16(k+1)$ for any $\beta \leq \frac{n}{4}$. Let X be the distortion of an edge chosen uniformly at random from $GRID_n$. X can only take on non-negative integer values, so the expected distortion is $E[X] = \sum_{i \geq 1} Prob(X \geq i)$. For $i \leq \frac{n}{64(k+1)}$, let $\beta = 16(k+1)i$. Then,

$$\begin{aligned}
E[X] &= \sum_{i \geq 1} Prob(X \geq i) \\
&> \sum_{i \geq 1}^{\lfloor n/64(k+1) \rfloor} \frac{n^2}{24 \cdot 16(k+1)i \cdot 2n(n-1)} \\
&> \sum_{i \geq 1}^{\lfloor n/64(k+1) \rfloor} \frac{1}{2 \cdot 24 \cdot 16(k+1)i} \\
&= \frac{1}{768(k+1)} \sum_{i \geq 1}^{\lfloor n/64(k+1) \rfloor} \frac{1}{i} \\
&= \Omega(\log n)
\end{aligned}$$

Since H was arbitrarily chosen, then by Yao's MiniMax principle if $GRID_n$ is embedded into a distribution over treewidth- k graphs, there must be an edge with expected distortion of $\Omega(\log n)$.

5 Conclusions

It is interesting to note that the inability of minor closed families to approximate all graphs well supports the conjecture [8] [9] that minor closed families of graphs (and therefore distributions over such families) can be embedded into ℓ_1 with constant distortion. Since Linial *et al.* [13] showed a lower bound of $\Omega(\log n)$ for embedding arbitrary graph metrics into ℓ_1 , the conjecture further implies that there must be families of graphs which cannot be embedded into distributions over excluded minor graphs with distortion less than $\Omega(\log n)$.

However, the particular inability of minor closed families to approximate other minor closed families also eliminates one potential approach to embedding these families into ℓ_1 : Gupta *et al.* [8] showed that although treewidth-2 graphs and treewidth-1 graphs are both embeddable into ℓ_1 with constant distortion, treewidth-2 graphs are not embeddable into distributions over treewidth-1 graphs with constant distortion. We have shown that a similar structure holds for all higher treewidths. Thus, an approach which attempts to repeatedly embed bounded treewidth graphs into (distributions over) graphs with lower treewidth will not work.

Acknowledgements: The authors would like to thank an anonymous reviewer for pointing out the aforementioned folklore result. We would also like to thank Elias Koutsoupias for valuable discussions, and Adam Meyerson and Shailesh Vaya for helpful comments on previous drafts.

References

1. N. Alon, R. Karp, D. Peleg, and D. West, "A Graph-Theoretic Game and its Application to the k-Server Problem", *SIAM J. Comput.*, 24:1 (1998), pp. 78-100.
2. Y. Bartal, "Probabilistic approximation of metric spaces and its algorithmic applications", In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, 1996, pp. 184-193.
3. Y. Bartal, "On approximating arbitrary metrics by tree metrics", In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, 1998, pp. 161-168.
4. B. Baker, "Approximation Algorithms for NP-complete Problems on Planar Graphs", *J. ACM*, 41 (1994), pp. 153-180.
5. H. Bodlaender, "A partial k-arboretum of graphs with bounded treewidth", *Theoretical Computer Science*, 209 (1998), pp. 1-45.
6. R. Diestel, "Graph Theory", Springer-Verlag, New York, 1997.
7. J. Fakcharoenphol, S. Rao, and K. Talwar, "A Tight Bound on Approximating Arbitrary Metrics by Tree Metrics", In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, 2003, pp. 448-455.
8. A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair, "Cuts, trees and ℓ_1 -embeddings.", In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, 1999, pp. 399-408.
9. C. Chekuri, A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair, "Embedding k -Outerplanar Graphs into ℓ_1 ", In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2003, pp. 527-536.
10. M. Imase and B. Waxman, "Dynamic Steiner Tree Problem.", *SIAM J. Discrete Math.*, 4 (1991), pp. 369-384.
11. P. Klein, S. Plotkin, and S. Rao, "Excluded Minors, Network Decomposition, and Multicommodity Flow", In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, 1993, pp. 682-690.
12. T. Kloks, "Treewidth: Computations and Approximations", *Lecture Notes in Computer Science*, Vol. 842, Springer-Verlag, Berlin, 1994.
13. N. Linial, E. London, and Y. Rabinovich, "The geometry of graphs and some of its algorithmic applications", *Combinatorica*, 15 (1995), pp. 215-245.
14. Y. Rabinovich, "On Average Distortion of Embedding Metrics into the Line and into ℓ_1 ", In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, 2003, pp. 456-462.
15. Y. Rabinovich and R. Raz, "Lower Bounds on the Distortion of Embedding Finite Metric Spaces in Graphs", *Discrete & Computational Geometry*, 19 (1998), pp. 79-94.
16. S. Rao, "Small distortion and volume preserving embeddings for Planar and Euclidean metrics", In *Proceedings of the 15th Annual Symposium on Computational Geometry*, 1999, pp. 300-306.