Light Spanners in Bounded Pathwidth Graphs

Michelangelo Grigni and Hao-Hsiang Hung

Dept. of Math & CS, Emory University, {*mic,hhung2*}*@mathcs.emory.edu*

Abstract. Given an edge-weighted graph G and $\epsilon > 0$, a $(1+\epsilon)$-spanner is a spanning subgraph G' whose shortest path distances approximate those of G within a factor of $1+\epsilon$. For G from certain graph families (such as bounded genus graphs and apex graphs), we know that *light* spanners exist. That is, we can compute a $(1+\epsilon)$-spanner G' with total edge weight at most a constant times the weight of a minimum spanning tree. This constant may depend on ϵ and the graph family, but not on the particular graph G nor on the edge weighting. The existence of light spanners is essential in the design of approximation schemes for the metric TSP (the traveling salesman problem) and similar graph-metric problems.

In this paper we make some progress towards the conjecture that light spanners exist for every minor-closed graph family: we show that light spanners exist for graphs with bounded pathwidth, and they are computed by a greedy algorithm. We do this via the intermediate construction of light *monotone* spanning trees in such graphs.

1 Introduction

1.1 Light Spanners

Suppose G is a connected undirected graph where each edge e has length (or weight) $w(e) \geq 0$. Let $d_G(u,v)$ denote the length of the shortest path between vertices u and v. Suppose G' is a spanning subgraph of G, where each edge of G' inherits its weight from G; evidently $d_G(u,v) \leq d_{G'}(u,v)$. Fix $\epsilon > 0$. If $d_{G'}(u,v) \leq (1+\epsilon) \cdot d_G(u,v)$ (for all u,v), then we say that G' is a $(1+\epsilon)$-*spanner* of G. In other words, the metric $d_{G'}$ closely approximates the metric d_G.

Let $w(G')$ denote the total edge weight of G', and let $\mathrm{MST}(G)$ denote the minimum weight of a spanning tree in G. We are interested in conditions on G that guarantee the existence of a $(1 + \epsilon)$-spanner G' with bounded $w(G')/\mathrm{MST}(G)$. Suppose $\mathcal{G}$ is a family of undirected graphs. We say $\mathcal{G}$ *has light spanners* if the following holds: for every $\epsilon > 0$ there is a bound $f(\epsilon)$, so that for any edge-weighted G from $\mathcal{G}$, G has a $(1 + \epsilon)$-spanner G' with $w(G') \leq f(\epsilon) \cdot \mathrm{MST}(G)$. Less formally, we say that G' is a *light spanner* for G. Note $f(\epsilon)$ depends on ϵ and $\mathcal{G}$, but not on G or w.

We know that if a graph family has unbounded clique minors, then it does not have light spanners; just consider a clique with uniform edge weights. We conjecture the converse [7]:

Conjecture 1. Any graph family with a forbidden minor has light spanners.

Our pursuit of this conjecture is guided by the Robertson-Seymour theory [12], which characterizes minor-closed graph families using four elements: bounded genus graphs, apices, vortices, and repeated clique-sums. We already know that if G has bounded genus or is an apex graph, then it has light spanners [7,8]. In particular vortices are bounded pathwidth subgraphs, stitched inside the faces of a bounded genus graph.

1.2 Motivation

Conjecture 1 seems like a natural question, and its proof would address the "main difficulty" discussed in the concluding remarks of Demaine *et al.* [6], in the general context of approximation algorithms on weighted graphs. As a specific motivating problem, we review some results on the *metric TSP*, the Traveling Salesman Problem with triangle inequality. (For some other problems, see [3,4].)

We are given an edge-weighted graph G, and we seek a cyclic order of its vertices with minimum total distance as measured by d_G. Equivalently, we want a minimum weight cyclic tour in G visiting each vertex at least once. Let $\text{OPT}(G)$ denote the minimum tour weight; it is well known that $\text{MST}(G) \leq \text{OPT}(G) \leq 2 \cdot \text{MST}(G)$. We seek an approximation scheme: an algorithm which takes as inputs the weighted graph G and $\epsilon > 0$, and which outputs a tour with weight at most $(1 + \epsilon) \cdot \text{OPT}(G)$.

The problem is MAX SNP-hard [11], so we consider approximation schemes where the input graph G is restricted to some graph family $\mathcal{G}$ (e.g., planar graphs). We would like a PTAS (an approximation scheme running in time $O(n^{g(\epsilon)})$, for some function g), or better yet an EPTAS (an approximation scheme running in time $O(g(\epsilon) \cdot n^c)$, where the constant c is independent of ϵ).

Suppose $\mathcal{G}$ is a graph family, and that for any $G \in \mathcal{G}$ we can compute a $(1+\epsilon)$-spanner G' with $w(G') \leq f(\epsilon) \cdot \text{MST}(G)$. Then we may attempt to design a PTAS (or an EPTAS) for the metric TSP on $\mathcal{G}$, as follows:

1. On input G and ϵ, first compute G', a $(1 + \epsilon/2)$-spanner of G, with weight at most $f(\epsilon/2) \cdot \text{MST}(G)$.
2. Choose $\delta = (\epsilon/2)/f(\epsilon/2)$. Apply some algorithm finding a tour in G' with cost at most $\text{OPT}(G') + \delta \cdot w(G')$.
3. Return the tour, with cost at most $(1+\epsilon/2)\cdot\text{OPT}(G)+\delta\cdot(f(\epsilon/2)\cdot\text{MST}(G)) \leq (1 + \epsilon) \cdot \text{OPT}(G)$. (For other metric optimization problems, it may be less trivial to lift a solution from G' back to G.)

Step 2 looks like the original problem, except now we allow an error term proportional to $w(G')$ instead of $\text{OPT}(G')$. This approach has already succeeded for planar graphs [2,9] and bounded genus graphs [6,7].

A recent result of Demaine *et al.* [5, Thm. 2] implies a PTAS for metric TSP when $\mathcal{G}$ is *any* graph class with a fixed forbidden minor. Since we do not know that $\mathcal{G}$ has light spanners, for step 1 they substitute a looser result [8], finding a $(1 + \epsilon)$-spanner G' with weight $O((\log n)/\epsilon) \cdot \text{MST}(G)$ (the hidden constant depending on $\mathcal{G}$). In step 2 their algorithm runs in time $2^{O(1/\delta + \log n)}$. Their $1/\delta$ is $O(w(G')/(\text{MST}(G) \cdot \epsilon)) = O((\log n)/\epsilon^2)$, so their running time is

$n^{O(1/\epsilon^2)}$. If we could compute light spanners for $\mathcal{G}$, then δ would improve to something independent of n, and this would yield an EPTAS for metric TSP on $\mathcal{G}$. (Or alternatively, it would yield an approximation scheme allowing ϵ to slowly approach zero, as long as $1/\delta$ stays $O(\log n)$.)

1.3 Our Work

In this paper we make some progress towards Conjecture 1: we show that light spanners exist for bounded pathwidth graphs.

Theorem 1. *Bounded pathwidth graphs have light spanners, computable by a greedy algorithm.*

We prove this in Section 3. This result is not algorithmically interesting by itself, since metric TSP (and many other problems) is exactly solvable in polynomial time when G has bounded pathwidth, or even bounded treewidth. Rather, we regard this as progress towards the conjecture, and towards an EPTAS for metric TSP (and similar problems) on graphs with forbidden minors. See Section 4 for some further remarks.

2 Preliminaries

2.1 Charging Schemes

In order to exhibit light spanners in a weight-independent way, we use *charging schemes* [8]. (We use the notion called "0-schemes" in [8], not the more general "ϵ-schemes" required for apex graphs.) Suppose each edge of graph G can hold some quantity of *charge*, initially zero. A *detour* is an edge $e \in E$ and a path P such that $e + P$ is a simple cycle in G. For each detour (e, P) we introduce a variable $x_{(e,P)} \geq 0$. Each $x_{(e,P)}$ describes a *charging move*: it subtracts $x_{(e,P)}$ units of charge from edge e, and adds $x_{(e,P)}$ units of charge to each edge of P. When $x_{(e,P)} > 0$, we say "e charges P".

Given graph G, a spanning tree T, and a number v, a *charging scheme from G to T of value v* is an assignment of nonnegative values to the $x_{(e,P)}$ variables (i.e., a fractional sum of detours) meeting the three conditions listed below. Here $\text{out}(e)$ denotes the total charge subtracted from edge e, $\text{in}(e)$ denotes the total charge added to e (as part of various detour paths), and $\text{net}(e) = \text{in}(e) - \text{out}(e)$ is the total charge on e after all the moves are done:

$$
\begin{aligned}
&(1) \quad \text{out}(e) \geq 1 \quad \text{for all } e \in G - T, \\
&(2) \quad \text{net}(e) \leq 0 \quad \text{for all } e \in G - T, \\
&(3) \quad \text{net}(e) \leq v \quad \text{for all } e \in T.
\end{aligned}
$$

Note "$e \in G - T$" means e is an edge of G but not T. As we'll see in Theorem 2, charging schemes imply light spanners.

Definition 1. *An* acyclic scheme *is a charging scheme with two additional conditions:*

(4) *If edge e charges some path, then $e \in G - T$.*

(5) *There is an ordering of the edges such that whenever edge e_1 charges a path containing edge e_2, e_1 precedes e_2.*

For example, planar graphs have integral acyclic schemes of value $v = 2$ [1].

Definition 2. *Suppose we have detours (e_1, P_1) and (e_2, P_2), with $e_2 \in P_1$ and $e_1 \notin P_2$. Their* shortcut *is the detour (e_1, P'), where P' is the path derived from P_1 by replacing e_2 with P_2, and then reducing that walk to a simple path.*

Lemma 1. *Suppose we have an acyclic scheme of value v from G to T, and an edge e in $G - T$. Then there is an acyclic scheme of value v from $G - e$ to T.*

Proof. Let $e_2 = e$. While in(e_2) is positive, we find some e_1 charging a path P_1 containing e_2. Since net$(e_2) \le 0$, e_2 also charges some path P_2. P_2 cannot contain e_1, since the scheme is acyclic. Let $\alpha = \min(x_{(e_1, P_1)}, x_{(e_2, P_2)})$. Now reduce both $x_{(e_1, P_1)}$ and $x_{(e_2, P_2)}$ by α, and increase $x_{(e_1, P')}$ (their shortcut) by α. After this change all the conditions are still satisfied, except possibly for condition (1) at e_2. Repeat until in(e_2) reaches zero. Finally remove e_2 and any remaining charges out of e_2. □

Theorem 2. *Suppose G is a graph with spanning tree T, and we have an acyclic scheme from G to T of value v. Then for any $\epsilon > 0$, and for any non-negative edge-weighting w on G, a simple greedy algorithm finds a $(1 + \epsilon)$-spanner G' in G containing T, with total weight $w(G') \le (1 + v/\epsilon) \cdot w(T)$.*

We use the following greedy algorithm of Althöfer *et al.* [1], modified to force the edges of T into G':

> Spanner(G, T, $1 + \epsilon$):
> $G' = T$
>   for each edge $e \in G - T$, in non-decreasing $w(e)$ order
>   if $(1 + \epsilon) \cdot w(e) < d_{G'}(e)$ then
>   add edge e to G'
>   return G'

The proof of Theorem 2 is a variant of previous arguments by LP duality [7,8], for completeness we sketch it here.

Proof. Since G' is computed by the greedy algorithm, it is clearly a $(1 + \epsilon)$-spanner of G containing T; the issue is to bound its weight $w(G')$. By Lemma 1 we have an acyclic scheme from G' to T of value v.

Consider a detour (e, P) in G' with $e \notin T$. We claim $(1 + \epsilon) \cdot w(e) < w(P)$ (to see this, compare e with the last edge inserted by the algorithm on the cycle $e + P$). Multiply through by $x_{(e,P)}$ and we have this:

$$x_{(e,P)} \cdot \epsilon \cdot w(e) \le x_{(e,P)} \cdot (w(P) - w(e))$$

When $e \in T$ this is still valid, since $x_{(e,P)} = 0$. Now sum over all detours (e, P):

$$\sum_{(e,P)} x_{(e,P)} \cdot \epsilon \cdot w(e) \leq \sum_{(e,P)} x_{(e,P)} \cdot (w(P) - w(e))$$

$$\epsilon \cdot \sum_{e \in G'} w(e) \cdot \text{out}(e) \leq \sum_{e \in G'} w(e) \cdot \text{net}(e)$$

$$\epsilon \cdot w(G' - T) \leq v \cdot w(T)$$

So $w(G') = w(T) + w(G' - T) \leq (1 + v/\epsilon) \cdot w(T)$. $\qquad\qquad\square$

2.2 Bounded Pathwidth and Monotone Trees

Suppose $G = (V, E)$ is a graph, P is a path (disjoint from G), and $\mathcal{B} = (B_i)_{i \in P}$ is a collection of subsets of V (bags) indexed by vertices i in P. We call the pair $(P, \mathcal{B})$ a *path decomposition* of G if the following conditions hold: (1) $\bigcup_{i \in P} B_i = V$; (2) for every edge $\{u, v\} \in E$, there is at least one bag B_i with $\{u, v\} \subseteq B_i$; (3) for every $v \in V$, $\{i : v \in B_i\}$ is connected (an interval) in P. The *pathwidth* of the decomposition is the maximum bag size minus one, and the pathwidth of G is the minimum pathwidth of any path decomposition of G.

Given $(P, \mathcal{B})$, we may lay out P on the line, and regard G as a subgraph of an interval graph. That is, for each vertex v we have a line interval I_v (corresponding to an interval in P), and we have $I_u \bigcap I_v \neq \emptyset$ whenever $\{u, v\} \in E$, and at most $k+1$ intervals overlap at any point of the line. For convenience we may eliminate ties via small perturbation, so that all the interval endpoints are distinct. In particular, let $\text{left}(v)$ denote the leftmost point of I_v.

Suppose T is a rooted tree in G. We say T is a *monotone tree* if for every vertex v in T with parent p, we have $\text{left}(p) < \text{left}(v)$. When T is a path rooted at an endpoint, we say it is a *monotone path*. In particular if T is a monotone spanning tree in G, then from any vertex v, we can find a monotone path in T from v to the root of T (the vertex with the leftmost interval). For this process, it is convenient to imagine that edges connect intervals at their leftmost intersection point.

3 Main Argument

We are given $\epsilon > 0$, a connected edge-weighted graph G with n vertices, and an interval representation $\{I_v\}$ of G with pathwidth k. We want to find a $(1 + \epsilon)$-spanner G' in G of low weight. First we apply some reductions to simplify G:

Nice Decomposition. We may assume that each pair of consecutive bags (as vertex sets) differ by only one vertex. This can be enforced by an argument similar to the construction of *nice* tree-decompositions [10]: if two consecutive bags differ on $m \geq 2$ vertices, we introduce $m - 1$ intermediate bags, in such a way that each pair differs on only one vertex, and we do not increase the maximum bag size. This does not modify G at all.

Bounded Degree Assumption. We may assume each vertex appears in $O(k)$ bags, and so the maxdegree of G is $O(k)$. To enforce this, we copy the bags of G from left to right. After each group of k original bags, ending with a bag B, we insert $|B|$ "replacer" bags, each of which replaces one vertex $v \in B$ with a copy v', connected to v by an edge of length zero. This ends with a bag B', where every vertex $v \in B$ has been replaced by a copy $v' \in B'$. See Figure 1. We continue in this way (using the copies in place of the originals) across the entire path decomposition. If we aren't careful the pathwidth may increase by one, but this does not matter for our asymptotic results. The original graph is obtained by contracting a set S of weight-zero edges in the modified graph. So given a spanner G' in this modified graph, we may contract S in $G' \cup S$ to recover a spanner (of no greater weight) in the original.

Completion Assumption. We may assume that G is *completed*; that is, it contains all edges allowed by its overlapping intervals. In other words: we have a clique in each bag, G is an interval graph. For each absent edge $e = \{u, v\}$, we simply add it with weight $w(e)$ equal to the shortest path length $d_G(u, v)$. This does not change d_G at all. Given a spanner G' in the completed graph, we recover a spanner in the original graph by replacing each completion edge by the corresponding shortest path.


Fig. 1. Each vertex v in bag B is replaced by v' in bag B'.

Proof (of Theorem 1). We assume all the above reductions have been applied: the input graph G is a connected edge-weighted interval graph of width k, each bag in its path decomposition introduces at most one vertex, and each vertex of G has degree $O(k)$.

By Lemma 2 (below), we compute a monotone spanning tree T with $w(T) = O(k^2) \cdot \mathrm{MST}(G)$. By Lemma 3 (below), we exhibit an acyclic charging scheme from G to T of value $v = O(k)$. Finally we apply the greedy algorithm, which computes a $(1 + \epsilon)$-spanner G'. By Theorem 2, $w(G') \leq (1 + v/\epsilon) \cdot w(T) = O(k^3/\epsilon) \cdot \mathrm{MST}(G)$. $\qed$

Lemma 2. *Given G as above, it contains a monotone spanning tree T with $w(T) \leq O(k^2) \cdot MST(G)$.*

Proof. Choose a minimum spanning tree T^*, so $w(T^*) = \text{MST}(G)$. Let I_l and I_r be the leftmost and rightmost intervals. Let P_1 be a shortest path from I_l to I_r; since G is completed, we may assume P_1 is monotone, as in Figure 2. Note $w(P_1) \leq w(T^*)$.

Consider the components $T_1^*, T_2^*, ..., T_m^*$ of $T^* - V(P_1)$. Let e_i be an edge connecting the leftmost point of T_i^* to a vertex of P_1 (it exists by completion). For each T_i^*, we recursively compute a monotone spanning tree T_i of $G[V(T_i^*)]$. Finally, $T = P_1 \cup \bigcup_i (T_i \cup e_i)$.

It is clear that T is monotone, but we must account for the total weight of $w(T)$. For each component T_i^*, let f_i be an edge of T^* connecting T_i^* to P_1 (there must be at least one). By triangle inequality, we see $w(e_i)$ is at most $w(T_i^*) + w(f_i) + w(P_{1,i})$, where $P_{1,i}$ is a subpath of P_1 from the endpoint of e_i to the endpoint of f_i. Note the f_i's and T_i^*'s are disjoint parts of T^*, but the subpaths may overlap inside P_1.


Fig. 2. P_1 and the T_i^* subtrees. Each f_i in T^* is replaced by an e_i in T.

An edge $e \in P_1$ appears in at most $k - 1$ of the $P_{1,i}$ subpaths, since each subpath witnesses another vertex (from T_i^*) that must appear in the bag with e. So $\sum_i w(e_i) \leq \sum_i [w(f_i) + w(T_i^*) + w(P_{1,i})] \leq w(T^*) + (k-1)w(P_1) \leq k \cdot w(T^*)$. Since $w(T) \leq O(k \cdot w(T^*)) + \sum_i w(T_i)$ and $\sum_i w(T_i^*) \leq w(T^*)$, a simple depth-$k$ recursion finishes our bound. □

Remark: we do not have to compute T as in Lemma 2; it suffices to use any light enough monotone spanning tree. A natural choice is to let T be the *lightest* monotone spanning tree, which we compute as follows. Start with just the root (in the leftmost bag), and grow the tree in a left-to-right scan of the bags: each time a bag B introduces a new vertex v, add an edge connecting v to its nearest neighbor in B (which is already in T).

In the completed G, a *triangle move* is a charging move where a non-tree edge e charges a path P of length two, where at most one edge of P is not in T. We now define $T^{(2)}$, a graph whose edges represent triangle moves. Each vertex jk of $T^{(2)}$ corresponds to an edge $\{j, k\}$ in G. We also represent the vertex jk by the interval $I_{jk} = I_j \cap I_k$. To define the edges of $T^{(2)}$, we first define a *parent* for each vertex jk. If $\{j, k\}$ is an edge of T, then jk has no parent. Otherwise, suppose $\mathrm{left}(j) < \mathrm{left}(k)$ (else swap them), and let i be the parent of k in T; i must exist since k is not the root. Note $\{i, j, k\}$ is a triangle in G. Now we say the parent of jk is ij, and we add the edge $\{ij, jk\}$ in $T^{(2)}$. Note $\mathrm{left}(ij) < \mathrm{left}(jk)$, so these parent links are acyclic. Thus $T^{(2)}$ is a forest, with each component rooted at a vertex corresponding to an edge of T. Figure 3 illustrates a simple monotone tree T and its forest $T^{(2)}$.


(a) a monotone tree T (b) $T^{(2)}$ produced from T

Fig. 3. Horizontal lines are intervals, dashed verticals are edges of T.

Lemma 3. *Given G as above with a monotone spanning tree T, there is an acyclic charging scheme from G to T of value $O(k)$.*

Proof. Recall $T^{(2)}$ is a forest. Fix a component C of $T^{(2)}$; it is a tree, rooted at a vertex r corresponding to an edge of T, and that is the only such vertex in C. Consider a directed Euler tour of C, traversing each edge twice. Delete each tour edge out of r, so we get a list of directed paths, each of the form

$$e_1 \to e_2 \to \cdots e_m \to r$$

where each vertex e_i corresponds to some edge of $G - T$. Since C is a tree, these paths are vertex disjoint (except at r). However, a vertex may appear more than once on the same path; call an appearance e_i a *repeat* if the same vertex appeared earlier on the path. Let $\mathcal{P}$ be the collection of all these paths, from all components of $T^{(2)}$.

We now propose a charging scheme (which fails to be acyclic). Recall how we constructed edges in $T^{(2)}$: we connect each vertex jk (corresponding to an
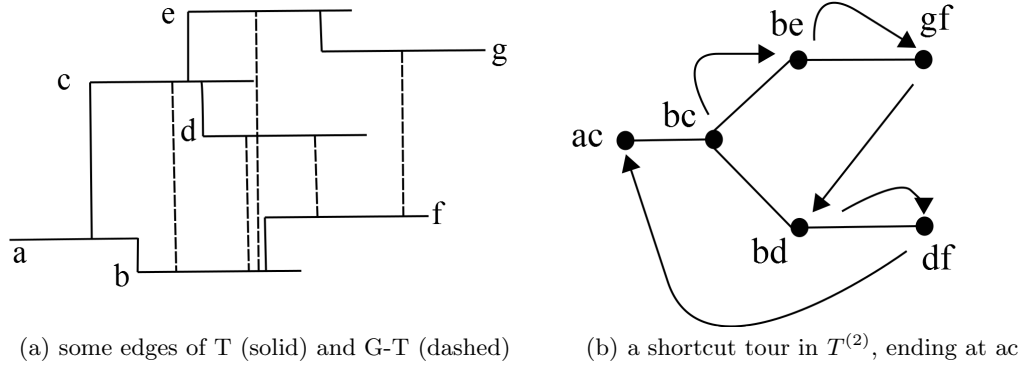
(a) some edges of T (solid) and G-T (dashed)   (b) a shortcut tour in $T^{(2)}$, ending at ac

Fig. 4. Edges of G (left) are vertices of $T^{(2)}$ (right)

edge of $G - T$) to its parent ij. If a path in $\mathcal{P}$ traverses this edge in the direction $jk \to ij$, we add the triangle move where edge $\{j, k\}$ charges one unit to path $j - i - k$. If a path traverses this edge in the other direction $ij \to jk$ (so ij is not a tree edge), we add the triangle move where edge $\{i, j\}$ charges one unit to path $i - k - j$. In either direction, the tree edge $\{i, k\}$ is charged.

For an edge $e \in G - T$, the corresponding vertex appears at least once on a path, and it has at least as many out-edges as in-edges, so our proposed scheme satisfies conditions (1) and (2). For an edge $e \in T$, we must bound the number of times it is charged. Since G has maxdegree $O(k)$, e appears in $O(k)$ distinct triangles, and it is charged at most twice per triangle (this includes the charges it receives in its role as r). So if we choose $v = O(k)$, condition (3) is satisfied. Also there are no charges out of tree edges, so condition (4) is also satisfied.

However, this charging scheme does not satisfy condition (5); if a vertex (corresponding to an edge $e \in G - T$) has a repeat appearance on its path, then there is no consistent way to order the edges. To fix this, we eliminate all "repeat" appearances using shortcuts. That is, whenever we have a sequence $e_1 \to e_2 \to e_3$ where e_2 is a repeat, we shortcut out e_2. Note such shortcuts can be combined. For example if we have a sequence $e_1 \to e_2 \to e_3 \to e_4 \to e_5$, corresponding to four triangle moves, it is possible to shortcut out e_2, e_3, e_4 (in any order), and the result is a single charge from e_1 to a path containing e_5 (the rest of the charged path is all tree edges). After eliminating all repeats by shortcuts, we get the desired acyclic scheme. $\qquad\square$

4 Conclusion and Further Work

Regarding our main result, it is not clear whether we really need to force the edges of a monotone T in the greedy spanner computation. Also, we might hope to reduce the $O(k^3)$ factor to something smaller.

The next obvious target is bounded treewidth graphs, a prerequisite for handling clique sums as in the Robertson-Seymour characterization.

There are several obvious directions to try extending the current approach to further minor-closed graph families. First, as extensions of Theorem 1, we propose two open problems: show light spanners for a planar graph with a single vortex, and show light spanners for a path-like clique-sum of planar graphs. For these cases it may help to compose multiple charging schemes into an "ϵ-scheme", as was necessary for apex graphs [8]. As usual, the main difficulty is that we have no control over the MST topology; if the MST has a nice topology (e.g. some form of monotonicity), then we would be done.

Given a bounded treewidth graph, we can still define the notion of a monotone spanning tree T. We choose roots in the decomposition tree and T; whenever vertex v has parent p in T, we require p to be in the bag containing v which is closest to the root. If we can find such a T that is light enough, then we could repeat the rest of our argument from the bounded pathwidth case. However, there is an obstacle: the light monotone tree might not exist.

Theorem 3. *There is an edge-weighted graph G with a bounded treewidth decomposition, such that any monotone spanning tree T in the completion of G has weight $\Omega(\lg n) \cdot MST(G)$.*


Fig. 5. A bounded treewidth graph with no light monotone tree. The solid path P is the minimum spanning tree. The horizontal edge in each leaf bag has weight one, all other solid edges of P have weight zero. All other edges (in particular, the dashed ones) have weight equal to the distance in P between its endpoints.

Proof. We construct G as follows (see Figure 5). We start with a balanced binary tree with n nodes, think of this as our tree of bags. We assign 3 nodes to each internal bag, and 2 to each leaf bag. We connect these vertices by a path P as shown; each edge of P in a leaf bag has weight one, all other edges in P have weight zero. (Note we must grow our bags a bit to support all these edges of P.)

Now in any monotone tree T for G, for each internal bag, the "bottom" vertex of the three must be connected to one of the other two (its parent in T); in other words, we must pick one of the dashed edges shown in each internal bag.

The main observation is that if we sum up the weights of these selected edges over one level of the decomposition tree, their total is already a constant fraction of $w(P)$. Summing over all levels, the total weight $w(T)$ is $\Omega(\log n) \cdot w(P)$, as claimed. □

References

1. Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9:81–100, January 1993.
2. Sanjeev Arora, Michelangelo Grigni, David R. Karger, Philip N. Klein, and Andrzej Woloszyn. A polynomial-time approximation scheme for weighted planar graph TSP. In *SODA*, pages 33–41, 1998.
3. André Berger, Artur Czumaj, Michelangelo Grigni, and Hairong Zhao. Approximation schemes for minimum 2-connected spanning subgraphs in weighted planar graphs. In *in Proc. of ESA*, pages 472–483, 2005.
4. André Berger and Michelangelo Grigni. Minimum weight 2-edge-connected spanning subgraphs in planar graphs. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *ICALP*, volume 4596 of *Lecture Notes in Computer Science*, pages 90–101. Springer, 2007.
5. Erik D. Demaine, MohammadTaghi Hajiaghayi, and Ken ichi Kawarabayashi. Contraction decomposition in h-minor-free graphs and algorithmic applications. In *Proceedings of the 43rd ACM Symposium on Theory of Computing (STOC 2011)*, page to appear, June 6–8 2011.
6. Erik D. Demaine, MohammadTaghi Hajiaghayi, and Bojan Mohar. Approximation algorithms via contraction decomposition. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '07, pages 278–287, Philadelphia, PA, USA, 2007. Society for Industrial and Applied Mathematics.
7. Michelangelo Grigni. Approximate TSP in graphs with forbidden minors. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, ICALP '00, pages 869–877, London, UK, 2000. Springer-Verlag.
8. Michelangelo Grigni and Papa Sissokho. Light spanners and approximate TSP in weighted graphs with forbidden minors. In *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '02, pages 852–857, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.
9. Philip N. Klein. A linear-time approximation scheme for tsp for planar weighted graphs. In *Proceedings, 46th IEEE Symposium on Foundations of Computer Science*, pages 146–155, 2005.

10. Ton Kloks. *Treewidth, Computations and Approximations*, volume 842 of *Lecture Notes in Computer Science*. Springer, 1994.
11. Christos H. Papadimitriou and M. Yannakakis. The Traveling Salesman Problem with distances one and two. *Mathematics of Operations Research*, 18:1–11, 1993.
12. Neil Robertson and Paul D. Seymour. Graph minors. XVI. Excluding a non-planar graph. *J. Comb. Theory, Ser. B*, 89(1):43–76, 2003.