Equivalence of Local Treewidth and Linear Local Treewidth and its Algorithmic Applications (extended abstract)

Erik D. Demaine and MohammadTaghi Hajiaghayi

Laboratory for Computer Science, Massachusetts Institute of Technology,
200 Technology Square, Cambridge, MA 02139, U.S.A.
{edemaine,hajiagha}@theory.lcs.mit.edu

Abstract. We solve an open problem posed by Eppstein in 1995 [16, 17] and re-enforced by Grohe [18, 19] concerning locally bounded treewidth in minor-closed families of graphs. A graph has *bounded local treewidth* if the subgraph induced by vertices within distance r of any vertex has treewidth bounded by a function of r (not n). Eppstein characterized minor-closed families of graphs with bounded local treewidth as precisely minor-closed families that minor-exclude an apex graph, where an *apex graph* has one vertex whose removal leaves a planar graph. In particular, Eppstein showed that all apex-minor-free graphs have bounded local treewidth, but his bound is doubly exponential in r, leaving open whether a tighter bound could be obtained. We improve this doubly exponential bound to a linear bound, which is optimal. In particular, any minor-closed graph family with bounded local treewidth has linear local treewidth. Our bound generalizes previously known linear bounds for special classes of graphs proved by several authors. As a consequence of our result, we obtain substantially faster polynomial-time approximation schemes for a broad class of problems in apex-minor-free graphs, improving the running time from $2^{2^{2^{O(1/\varepsilon)}}} n^{O(1)}$ to $2^{O(1/\varepsilon)} n^{O(1)}$.

1 Introduction

Many problems are inapproximable beyond $\Theta(\log n)$ or $\Theta(n^{\varepsilon})$ factors for general graphs, yet for classes of graphs with additional structural properties, we can obtain polynomial-time approximation schemes (PTAS). The first collection of results along these lines was for planar graphs [23]. These results were later generalized to H-minor-free graphs for any fixed graph H [2]. Both of these results are impractical; for example, just to achieve an approximation ratio of 2, the base case of the planar-graph approach requires exhaustive solution of graphs of up to $2^{2^{400}}$ vertices [11]. To address this impracticality, Baker [5] introduced a practical approach for planar graphs based on decomposition into overlapping subgraphs of bounded outerplanarity. Specifically, Baker's approach obtains $(1+\varepsilon)$-approximation algorithms with running times of $2^{O(1/\varepsilon)} n^{O(1)}$ for many problems on planar graphs, such as maximum independent set, minimum dominating set, and minimum vertex cover. Chen [10] later generalized Baker's approach to obtain PTASs for $K_{3,3}$-minor-free graphs and K_5-minor-free graphs, but only for maximization problems.

Eppstein [17, 16] further generalized Baker's approach by replacing local regions of bounded outerplanarity with local regions of bounded treewidth. The resulting approximation algorithms apply to any graph of bounded local treewidth, a notion newly introduced by Eppstein [16, 17]. A graph has *bounded local treewidth* if the treewidth of the subgraph induced by the set of vertices at distance at most r from any vertex is bounded above by some function $f(r)$ independent of n. Eppstein also characterized all minor-closed families of graphs that have bounded local treewidth, showing that they are precisely apex-minor-free graphs, where an *apex graph* has a vertex whose removal leaves a planar graph. For example, $K_{3,3}$ and K_5 are apex graphs, and therefore apex-minor-free graphs, or equivalently graphs of bounded local treewidth, include planar graphs as a special case.

Eppstein's approach returns to the realm of impracticality, because his bound on local treewidth in a general apex-minor-free graph is doubly exponential in r: $2^{2^{O(r)}}$. Eppstein [17] improved this bound to linear—$O(r)$—for the special case of bounded-genus graphs, which are apex-minor-free [16]. This linear bound on local treewidth was later established for the special cases of $K_{3,3}$-minor-free and K_5-minor-free graphs [21], and then for single-crossing-minor-free graphs [14, 13]. A graph is *single-crossing* if it can be embedded in the plane with at most one crossing; thus, every single-crossing graph is an apex graph. Grohe [19] also established a linear bound on local treewidth for "apex-free almost-embeddable graphs"; see Section 2.4 for a definition.

In this paper, we prove that every apex-minor-free graph has linear local treewidth, generalizing all results of the previous paragraph. This result solves an open problem posed by Eppstein [17, 16] and mentioned in [14, 18, 19]. As a consequence, we obtain the surprising result that every minor-closed family of graphs with bounded local treewidth in fact has linear local treewidth. Along the way, we reprove Eppstein's characterization of minor-closed families of graphs with bounded local treewidth.

We recommend reading Section 4, which gives the intuition and high-level overview of the (difficult) proof of this result and the techniques involved. This intuition is intended to make sense (at a high level) even while skipping over the definitions in Sections 2 and 3.

Using our combinatorial results, we obtain substantially faster PTASs, improving the running time from $2^{2^{2^{O(1/\varepsilon)}}} n^{O(1)}$ to $2^{O(1/\varepsilon)} n^{O(1)}$, for many problems including hereditary maximization problems (e.g., maximum independent set), maximum triangle matching, maximum H-matching, maximum tile salvage, minimum vertex cover, minimum dominating set, minimum edge-dominating set, and subgraph isomorphism for a fixed pattern. We also substantially improve the running time of a key step in the general framework of [18] for deciding first-order properties about graphs of bounded local treewidth.

Grohe [19] developed PTASs for a few problems in the list above for graphs excluding a fixed minor H (not just when H is an apex graph). The running time of these algorithms includes a factor of $n^{f(H)}$ for a function f. In contrast, the running times of our algorithms replace this factor with $n^{3+\delta}$ for any $\delta > 0$, thus separating the dependence on n and H. (Both algorithms have another factor of the form $g(H, \varepsilon)$ where $1 + \varepsilon$ is the approximation ratio.)

Finally, it is worth metioning that recently approximation algorithms for H-minor-free graphs for a fixed graph H have been studied extensively; see e.g. [8, 20, 9, 22, 24]. In particular, it is generally believed that several algorithms for planar graphs can be generalized to H-minor-free graphs for any fixed H [20, 22, 24]. In fact, the main reason is that H-minor-free graphs are very general; the deep Graph-Minor Theory of Robertson and Seymour shows that any graph class that is closed under minors is characterized by excluding a finite set of minors.

This paper is organized as follows. Sections 2 and 3 introduce the terminology and basic concepts used throughout the paper. In particular, Section 3 introduces the necessary concepts and some of the main results from the Robertson and Seymour Graph Minor Theory. Section 4 provides a high-level description of our combinatorial results and their proof. Section 5 describes the many algorithmic applications of this combinatorial result, which use additional nonstandard tricks to improve running time. Section 6 establishes several structural results about graphs with bounded local treewidth and linear local treewidth. Section 7 gives a more-in-depth description of the proof of our combinatorial results. The full details of some of our lemmas share several components with the Graph Minor Theory, which is a long series of long papers. Therefore, we limit ourselves in this extended abstract to an intuitive description of the proofs of a few lemmas that cannot be directly obtained from the Graph Minor Theorems. Finally, we conclude with some remarks and open problems in Section 8.

2 Background

2.1 Preliminaries

All the graphs in this paper are undirected without loops or multiple edges. The reader is referred to standard references for appropriate background [7].

Our graph terminology is as follows. A graph G is represented by $G = (V, E)$, where V (or $V(G)$) is the set of vertices and E (or $E(G)$) is the set of edges. We denote an edge e between u and v by $\{u, v\}$. We define n to be the number of vertices of a graph when this is clear from context.

The *(disjoint) union* of two disjoint graphs G_1 and G_2, $G_1 \cup G_2$, is the graph G with merged vertex and edge sets: $V(G) = V(G_1) \cup V(G_2)$ and $E(G) = E(G_1) \cup E(G_2)$. We define the *$r$-neighborhood* of a set $S \subseteq V(G)$, denoted by $N_G^r(S)$, to be the set of vertices at distance at most r from at least one vertex of S; if $S = r$ we simply use the notation $N_G^r(v)$. For the ease of notation, we often use $N_G^r(v)$ instead of $G[N_G^r(v)]$ when it is clear from the context.

One way of describing classes of graphs is by using *minors*, introduced as follows. *Contracting* an edge $e = \{u, v\}$ is the operation of replacing both u and v by a single vertex w whose neighbors are all vertices that were neighbors of u or v, except u and v themselves. A graph G is a *minor* of a graph H if H can be obtained from a subgraph of G by contracting edges. A graph class $\mathcal{C}$ is a *minor-closed* class if any minor of any graph in $\mathcal{C}$ is also a member of $\mathcal{C}$. A minor-closed graph class $\mathcal{C}$ is *H-minor-free* if $H \notin \mathcal{C}$. For example, a planar graph is a graph excluding both $K_{3,3}$ and K_5 as minors.

2.2 Treewidth and local treewidth

The notion of treewidth was introduced by Robertson and Seymour [37] and plays an important role in their fundamental work on graph minors. To define this notion, first we consider the representation of a graph as a tree called *tree decomposition*. More precisely, a tree decomposition of a graph $G = (V, E)$, denoted by $TD(G)$, is a pair (χ, T) in which $T = (I, F)$ is a tree and $\chi = \{\chi_i | i \in I\}$ is a family of subsets of $V(G)$ such that: (1) $\bigcup_{i \in I} \chi_i = V$; (2) for each edge $e = \{u, v\} \in E$ there exists an $i \in I$ such that both u and v belong to χ_i; and (3) for all $v \in V$, the set of nodes $\{i \in I | v \in \chi_i\}$ forms a connected subtree of T. To distinguish between vertices of the original graph G and vertices of T in $TD(G)$, we call vertices of T *nodes* and their corresponding χ_i's *bags*. The maximum size of a bag in $TD(G)$ minus one is called the *width* of the tree decomposition. The *treewidth* of a graph G ($\mathrm{tw}(G)$) is the minimum width over all possible tree decompositions of G.

For many NP-complete problems, we can derive polynomial-time algorithms restricted to graphs of bounded treewidth using a general dynamic-programming approach similar to that on trees [4]. However, still the class of graphs of bounded treewidth is of limited size; we would like to solve NP-complete problems for wider classes of graphs. As mentioned before, Baker [5] developed several approximation algorithms to solve NP-complete problems for planar graphs. To extend these algorithms to other graph families, Eppstein [17] introduced the notion of bounded local treewidth, defined formally below, which is a generalization of the notion of treewidth. Intuitively, a graph has bounded local treewidth (or locally bounded treewidth) if the treewidth of an r-neighborhood of each vertex $v \in V(G)$ is a function of r, $r \in \mathbb{N}$, and not $|V(G)|$.

Definition 1. *The* local treewidth *of a graph G is the function $\mathrm{ltw}^G : \mathbb{N} \to \mathbb{N}$ that associates with every $r \in \mathbb{N}$ the maximum treewidth of an r-neighborhood in G. We set $\mathrm{ltw}^G(r) = \max_{v \in V(G)}\{\mathrm{tw}(G[N_G^r(v)])\}$, and we say that a graph class $\mathcal{C}$ has* bounded local treewidth (or locally bounded treewidth) *when there is a function $f : \mathbb{N} \to \mathbb{N}$ such that for all $G \in \mathcal{C}$ and $r \in \mathbb{N}$, $\mathrm{ltw}^G(r) \leq f(r)$. A class $\mathcal{C}$ has* linear local treewidth *if there are constants $c, d \in \mathbb{R}$ such that $\mathrm{ltw}^G(r) \leq cr + d$ for all $G \in \mathcal{C}$, $r \in \mathbb{N}$.*

Eppstein [17] showed that a minor-closed graph class $\mathcal{E}$ has bounded local treewidth if and only if $\mathcal{E}$ is H-minor free for some apex graph H.

2.3 Clique Sums

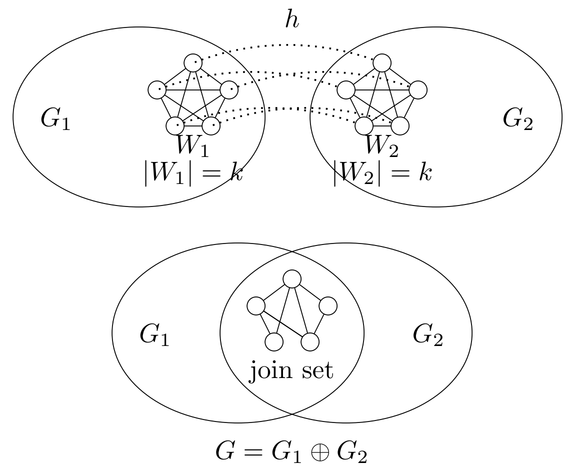

Fig. 1. Example of 5-sum of two graphs.

Suppose G_1 and G_2 are graphs with disjoint vertex-sets and $k \geq 0$ is an integer. For $i = 1, 2$, let $W_i \subseteq V(G_i)$ form a clique of size k and let G'_i ($i = 1, 2$) be obtained from G_i by deleting some (possibly no) edges from $G_i[W_i]$ with both endpoints in W_i. Consider a bijection $h : W_1 \to W_2$. We define a *k-sum* G of G_1 and G_2, denoted by $G = G_1 \oplus_k G_2$ or simply by $G = G_1 \oplus G_2$, to be the graph obtained from the union of G'_1 and G'_2 by identifying w with $h(w)$ for all $w \in W_1$. The images of the vertices of W_1 and W_2 in $G_1 \oplus_k G_2$ form the *join set*.

In the rest of this section, when we refer to a vertex v of G in G_1 or G_2, we mean the corresponding vertex of v in G_1 or G_2 (or both). It is worth mentioning that $\oplus$ is not a well-defined operator and it can have a set of possible results. See Figure 1 for an example of a 5-sum operation.

The following lemma whose intuition will play an important role in our results shows how the treewidth changes when we apply a clique-sum operation .

Lemma 1 ([21]). *For any two graphs G and H,* $\mathrm{tw}(G \oplus H) \leq \max\{\mathrm{tw}(G), \mathrm{tw}(H)\}$.

2.4 Clique-sum decompositions of H-minor-free graphs

Our result uses the deep theorem of Robertson and Seymour on graphs excluding a non-planar graph as a minor. Intuitively, Robertson-Seymour's theorem says for every graph H, every H-minor-free graph can be expressed as a tree-structure of "pieces", where each piece either has bounded size or is a graph which can be drawn in a surface in which H cannot be drawn, except for a bounded number of "apex" vertices and a bounded number of "local areas of non-planarity" called *vortices*. Here the bounds only depend on H.

Roughly speaking we say a graph G is *h-almost embeddable* in a surface S if there exists a set X of size at most h of vertices, called *apex vertices* or *apices*, such that $G - X$ can be obtained from a graph G_0 embedded in S by attaching at most h graphs of pathwidth at most h to G_0 along the boundary cycles $C_1, \cdots, C_h$ in an orderly way. More precisely:

Definition 2. *A graph G is h-*almost embeddable *in S if there exists a vertex set X of size at most h called* apices *such that $G - X$ can be written as $G_0 \cup G_1 \cup \cdots \cup G_h$, where*

- *G_0 has an embedding in S;*
- *the graphs G_i, called* vortices*, are pairwise disjoint;*
- *there are faces $F_1, \ldots, F_h$ of G_0 in S, and there are pairwise disjoint disks $D_1, \ldots, D_h$ in S, such that for $i = 1, \ldots, h$, $D_i \subset F_i$ and $U_i := V(G_0) \cap V(G_i) = V(G_0) \cap D_i$; and*
- *the graph G_i has a path decomposition $(\mathcal{B}_u)_{u \in U_i}$ of width less than h, such that $u \in \mathcal{B}_u$ for all $u \in U_i$. We note that the sets $\mathcal{B}_u$ are ordered by the ordering of their indices u as points in C_i, where C_i is the boundary cycle of F_i in G_0.*

An h-almost embeddable graph is called apex-free *if the set X of apices is empty.*

Now, the deep result of Robertson and Seymour is as follows.

Theorem 1 ([38]). *For every graph H there exists an integer $h \geq 0$ only depending on $|V(H)|$ such that every H-minor-free graph can be obtained by at most h-sums of graphs of size at most h and h-almost-embeddable graphs in some surfaces in which H cannot be embedded.*

In particular, if H is fixed, any surface in which H cannot be embedded has bounded genus. Thus, the summands in the theorem are h-almost-embeddable graphs in bounded-genus surfaces (In the rest of the paper, we consider bounded size pieces as almost-embeddable graphs whose bound genus parts are empty).

Unfortunately, since Theorem 1 is very general and has appeared in print very recently, not many other applications are known. However, this structural theorem plays an important role in obtaining the rest of the results of this paper, and we believe that it can be further useful in obtaining algorithms and proving theorems on graphs excluding a fixed graph H as a minor.

3 Technical Definitions

In this subsection we remind a very limited part of the machinery developed in Graph Minor papers that is used in our proofs.

A *surface* Σ is a compact 2-manifold, without boundary. A *line* in Σ is subset homeomorphic to $[0, 1]$. An O-arc is a subset of Σ homeomorphic to a circle. Let G be a graph 2-cell embedded in Σ. To simplify notations we do not distinguish between a vertex of G and the point of Σ used in the drawing to represent the vertex or between an edge and the line representing it. We also consider G as the union of the points corresponding to its vertices and edges. That way, a subgraph H of G can be seen as a graph H where $H \subseteq G$. We call by *region* of G any connected component of $\Sigma - E(G) - V(G)$. (Every region is an open set.) We use the notation $V(G), E(G)$, and $R(G)$ for the set of the vertices, edges and regions of G.

If $\Delta \subseteq \Sigma$, then $\overline{\Delta}$ denotes the *closure* of Δ, and the boundary of Δ is $\mathbf{bd}(\Delta) = \overline{\Delta} \cap \overline{\Sigma - \Delta}$. An edge e (a vertex v) is incident with a region r if $e \subseteq \mathbf{bd}(r)$ ($v \subseteq \mathbf{bd}(r)$).

A subset of Σ meeting the drawing only in vertices of G is called *G-normal*. If an O-arc is G-normal then we call it *noose*. The length of a noose is the number of its vertices. $\Delta \subseteq \Sigma$ is an open disc if it is homeomorphic to $\{(x, y) : x^2 + y^2 < 1\}$. We say that a disc D is *bounded* by a noose N if $N = \mathbf{bd}(D)$. A graph G 2-cell embedded in a connected surface Σ is *θ-representative* if every noose of length $< \theta$ is contractable (null-homotopic in Σ).

A *separation* of a graph G is a pair (A, B) of subgraphs with $A \cup B = G$ and $E(A \cap B) = \emptyset$, and its order is $|V(A \cap B)|$. Tangles were introduced by Robertson & Seymour in [28]. A *tangle of order* $\theta \geq 1$ is a set $\mathcal{T}$ of separations of G, each of order θ, such that

(i) for every separation (A, B) of G of order $< \theta$, $\mathcal{T}$ contains one of (A, B), (B, A)
(ii) if $(A_1, B_1), (A_2, B_2), (A_3, B_3) \in \mathcal{T}$ then $A_1 \cup A_2 \cup A_3 \neq G$.
(iii) if $(A, B) \in \mathcal{T}$ then $V(A) \neq V(G)$.

Let G be a graph embedded in a connected surface Σ. A tangle $\mathcal{T}$ of order θ is *respectful* if for every noose N in Σ with $|N \cap V(G)| < \theta$, there is a closed disc $\Delta \subseteq \Sigma$ with $\mathbf{bd}(\Delta) = N$ such that separation

$$(G \cap \Delta, G \cap \overline{\Sigma - \Delta}) \in \mathcal{T}.$$

Our proofs are based on the following results from Graph Minors papers by Robertson & Seymour.

Theorem 2 ((4.3) in [28]). *Let G be a graph with at least one edge. Then there is a tangle in G of order θ if and only if G has branch-width $\geq \theta$.*

Theorem 3 ((4.1) in [30]). *Let Σ be a connected surface, not a sphere, let $\theta \geq 1$, and let G be a θ-representative graph 2-cell embedded in Σ. Then there is a unique respectful tangle in G of order θ.*

Also in our proofs we use the notion of radial graph. Informally, the radial graph of a 2-cell embedded in Σ graph G is the bipartite graph R_G obtained by selecting a point in every region r of G and connecting it to every vertex of G incident to that region. However, a region maybe "incident more than once" with the same vertex, so one needs a more formal definition. A *radial drawing* R_G is a radial graph of a 2-cell embedded in Σ graph G if

1. $E(G) \cap E(R_G) = V(G) \subseteq V(R_G)$;
2. Each region $r \in R(G)$ contains a unique vertex $v_r \in V(R_G)$;
3. R_G is bipartite with a bipartition $(V(G), \{v_r : r \in R(G)\})$;
4. If e, f are edges of R_G with the same ends $v \in V(G)$, $v_r \in V(R_G)$, then $e \cup f$ does not bound a closed disc in $r \cup \{v\}$;
5. R_G is maximal subject to 1,2,3 and 4.

Finally, let $A(R_G)$ be the set of vertices, edges, and regions (collectively, *atoms*) in the radial graph R_G. According to Section 9 of [30] (see also [31]), the existence of a respectful tangle makes it possible to define a metric d on $A(R_G)$ as follows:

1. If $a = b$, then $d(a, b) = 0$.
2. If $a \neq b$, and a and b are interior to a contractible closed walk of radial graph of length $< 2\theta$, then $d(a, b)$ is half the minimum length of such a walk (here by *interior* we mean the direction in which the walk can be contracted).
3. Otherwise, $d(a, b) = \theta$.

We use these metrics very often in our proofs. Interestingly, this metric will be related to regular distance metric on grids, the property that we use in Lemma 13.

4 Local Treewidth of Apex-Minor-Free Graphs

The main result of this paper which has many algorithmic applications is as follows:

Theorem 4. *Any apex-minor-free graph has linear local treewidth.*

To obtain this result, we use the general structure of H-minor-free graphs given in Theorem 1. Also, we need the following theorem of Hajiaghayi et al. [21]:

Theorem 5 ([21]). *If G_1 and G_2 are graphs where $\mathrm{ltw}^{G_i}(r) \leq f(r)$, $f(r) \geq 0$ for all $r \in \mathbb{N}$, and $G = G_1 \oplus_k G_2$, then $\mathrm{ltw}^{G}(r) \leq f(r)$.*

Using Theorem 1 and Theorem 5, one can easily observe that we need to only prove the following theorem:

Theorem 6. *Each h-almost-embeddable graph in the clique-sum decomposition of an apex-minor-free graph G has linear local treewidth.*

The proof of Theorem 6 is lengthy and we delay the proof to Section 7. In the rest of this section, we mention some major ideas of the proof.

Our proof is based on a series of reductions, each of which uses the deep Graph Minor Theory of Robertson and Seymour. Each reduction converts a given graph into a "simpler" graph that has linear local treewidth if and only if the original graph has linear local treewidth. To achieve this equivalence, we must preserve the distances between pairs of vertices up to constant factors, thus roughly preserving the neighborhoods, and we must preserve the treewidth of these neighborhoods up to constant factors.

The first reduction effectively removes the vortices from a given almost-embeddable graph. A similar technique has been used by others, e.g., [19, 33]. Next we would like to use the property that the graph is apex-minor-free; however, only the original graph is apex-minor-free, and during the clique-sum decomposition, we may have introduced extra edges when the join set was completed into a clique. We call such edges *virtual edges*, and all other edges *actual edges*. One difficulty of Theorem 1 is that it does not guarantee that the virtual edges can be obtained by taking a minor of the original graph, and therefore the pieces may not be apex-minor-free. The second reduction overcomes this difficulty by obtaining some virtual edges by taking minors of the original graph, and removes other virtual edges which cannot be obtained, while still preserving linear local treewidth. We call the resulting graph an *approximation graph* of the original graph.

The next steps of the proof exploit the property that the approximation graph is apex-minor-free. Intuitively, apexes in the approximation graph play the role of the apex in the minor-excluded apex graph, and therefore the bounded-genus part of the approximation graph must (roughly) minor-exclude the planar part of the apex graph. For this property to be useful, the bounded-genus part of the approximation graph must be "spread out enough" in the sense of having high representativity.

If the bounded-genus graph has low representativity and is not planar, we can afford to decrease its genus by removing a small noose from the bounded-genus graph and placing all its vertices into the set of apexes. If the bounded-genus graph has high representativity, we use the resulting distance metric and the apex-minor-free property to decompose the graph into a clique-sum of several almost-embeddable graphs, re-introducing vortices in a special way. The bounded-genus parts of these almost-embeddable graphs are all planar and have no vortices *except* for the last which may not be planar and has a bounded number of vortices. This decomposition would seem to be making negative progress; however, the new apexes of the last graph have the special property that all their neighbors are within the vortices. In one shot, it is fairly easy to show that the last graph has linear local treewidth.

Finally, we deal with almost-embeddable graphs where the bounded-genus part is in fact planar (and has no vortices). We use the fact that a planar graph with treewidth w has as a minor an $O(w) \times O(w)$ grid. In contrast, Eppstein's approach considers general graphs instead of planar graphs, in which case we can only guarantee an $O(\log^{1/10} w) \times O(\log^{1/10} w)$ grid in a graph with treewidth w. This is one place where our bounds are stronger than Eppstein's, allowing us to ultimately get a linear bound $(O(r))$ on local treewidth instead of $2^{2^{O(r)}}$.

5 Algorithmic Applications: PTAS Results

In this section, we mention some algorithmic applications of Theorem 4. First, we start by hereditary maximization problem, which determine a property that if valid for an input graph is also valid for any induced subgraph of the input. For a property π, the maximum (weight) induced subgraph problem $MISP(\pi)$ is finding a maximum (weight) induced subgraph with the property. For example, we can search for an induced subgraph of maximum size that is chordal, acyclic, without cycles of a specified length, with edges, of maximum degree $d \geq 1$, bipartite or as clique [40]. According to Yannakakis almost all of these problems are NP-complete.

Theorem 7. *Let G be a non-negative vertex-weighted H-minor-free graph for an apex graph H and let $k \geq 1$ be an integer. The maximization problem MISP(π) for a hereditary property π over G has a PTAS of ratio $1 + 1/k$ of the optimal with worst-case running time in $O(k2^{3.698(ck+d)}|V|^{3+\varepsilon} + kTime_\pi 11/3(ck+d), |V|))$, where $Time_\pi(w, n)$ is the worst-case running time of WMISP(π) over an n-vertex partial w-tree whose tree decomposition is given. $Time_\pi(w, n)$ is nondecreasing as n increases and c and d are some constants which only depend on $|V(H)|$*

Proof. First we decompose graph G into several induced subgraphs, each of which having bounded treewidth, and mention some properties of these induced subgraphs. For $1 \leq i \leq k$ and $j \geq 0$, we define $L_{ij} = L[(j-1)k + i, jk + i - 2]$, where $L[e, f]$ contains all vertices of layers (a.k.a level or distance) e to f in a BFS tree from an arbitrary vertex $v \in V(G)$. Here we assume a layer is empty when its level number is not between zero and the total number of layers, e.g. consider $j = 0$. We note that there is no edge between L_{ij} and $L_{i(j+1)}$. Let $\mathcal{L}_i = \bigcup_{j \geq 0} L_{ij}$ and $G_i = G[\mathcal{L}_i]$. Here every vertex appears in exactly $k-1$ of the $\mathcal{L}_i$'s or G_i's (vertices in layer L_h only do not appear in $\mathcal{L}_i$ where i is congruent to $h + 1$ mod k). We label this fact by **[Fact a]**.

Then, we construct a tree decomposition of width $\frac{11}{3}(ck + d)$ for each $G[L_{ij}]$ using an algorithm of time $O(2^{3.698(ck+d)}n^{3+\varepsilon})$ of Amir [3]. Notice that here we use an approximation algorithms whose running time is so better that the exact algorithms such as Bodlaenders [6] used by others (this step is essential to get the claimed running time.) Since $G_i = \bigcup_{j \geq 0} G[L_{ij}]$, a tree decomposition of width $3(k-1) + 4$ for G_i can be constructed by gluing tree decompositions of $G[L_{ij}]$'s together (adding edges to become one tree) in $O(|V|)$ time (note that $G[L_{ij}]$'s are disjoint).

Next, we solve the MISP(π) on each G_i, $1 \le i \le k$. Since $|V(G_i)| \le |V(G)|$, Opt_i, the maximum weighted solution of MISP(π) over G_i, can be constructed in $Time_\pi(11/3(ck+d), |V(G)|)$.

Finally, we take Opt_m the solution with maximum weight among $Opt_1, Opt_2, \cdots, Opt_k$ as our solution for graph G, and show that it has a ratio $1 + 1/k$ of the optimal. Suppose Opt is the maximum weighted solution on graph G. We prove $\frac{weight(Opt)}{weight(Opt_m)} \le \frac{k}{k-1}$. Because of the hereditary property of MISP(π), we have:

$$weight(Opt \cap \mathcal{L}_i) \le weight(Opt_i) \tag{1}$$

Using 1, we have:

$$k \cdot weight(Opt_m) \ge \sum_{i=1}^{k} weight(Opt_i) \ge^{(1)} \sum_{i=1}^{k} weight(Opt \cap \mathcal{L}_i) =^{[Fact\ a]} (k-1) \cdot weight(Opt).$$

The claimed running time follows immediately from the running time of constructing the tree decomposition and solving MISP(π) for each G_i, and the number of G_i's. □

For example, for G be a non-negative vertex-weighted apex-minor-free graph, the maximum independent set problem admits a polynomial-time approximation scheme of ration $1+1/k$ with running time $O(k2^{3.698(ck+d)}|V|^{3+\varepsilon}+k4^{11/3(ck+d)}|V|)$ (see the dynamic programming for this problem in [4].)

In fact, the proof of Theorem 7 can be easily generalized for many other problems. For the sake of similarity of the proof, we only mention the general theorem.

Theorem 8. *Given an H-minor-free graph G, where H is an apex graph, there are PTASs with approximation ratio $1 + 1/k$ (or $1 + 2/k$) running in $O(c^k n)$ time (c is a constant) on graph G for hereditary maximization problems such as maximum independent set and other problems such as maximum triangle matching, maximum H-matching, maximum tile salvage, minimum vertex cover, minimum dominating set, minimum edge-dominating set, and subgraph isomorphism for a fixed pattern.*

Another algorithmic consequence of our results is in the context of fixed-parameter algorithms. Grohe [18] gives a general framework for deciding first-order properties of graphs of bounded local treewidth. Our result implies that these graphs in fact have linear local treewidth, substantially improving the worst-case bound on the running time of a key step in this framework. This improvement fixes one (but not all) of the bottlenecks of the algorithm mentioned by Grohe [18].

6 Linear Local Treewidth: Basic Properties

In this section we start demonstrating some basic properties of local treewidth, by which we are prepared to state the proof of Theorem 6 in the next section.

First we consider the relation between local treewidth of a graph before and after adding an edge between two non-adjacent vertices.

Lemma 2. *Let G be a graph with $\mathrm{ltw}^G(r) \le f(r)$ for all $r \in \mathbb{N}$. Also let G' be a graph obtained by adding an edge $\{u, w\}$ between two non-adjacent vertices u and v. Then $\mathrm{ltw}^{G'}(r) \le f(3r) + 1$.*

Proof. To show $\mathrm{ltw}^{G'}(r) \le f(3r) + 1$, we prove for any $v \in V(G)$ and for all $r \ge 0$, $\mathrm{tw}(N^r_{G'}(v)) \le f(3r)+1$. Since $f(r) \ge 0$, the claim is clear for $r = 0$. Thus we assume $r > 0$ in the rest of the proof. If $N^r_{G'}(v)$ does not contain the edge $\{u, w\}$, $N^r_{G'}(v) = N^r_G(v)$ and thus $\mathrm{tw}(N^r_{G'}(v)) \le f(r) \le f(3r)+1$ as desired. We now assume $N^r_{G'}(v)$ contains the edge $\{u, w\}$. Without loss of generality we assume that the distance of u from v called r_u is less than or equal to the distance of w from v called r_w. Notice that $N^r_{G'}(v) \subseteq N^r_G(v) \cup N^{r-r_u-1}_G(w)$. If $N^r_G(v)$ does not intersect $N^{r-r_u-1}_G(w)$ we have

$$tw(N^r_{G'}(v)) \le \max\{tw(N^r_G(v)), tw(N^{r-r_u-1}_G(w))\}+1 \le \max\{f(r), f(r-r_u-1)\}+1 \le f(r)+1 \le f(3r)+1.$$

Here we add one since we add u to all bags.

In the other hand if $N_G^r(v)$ intersects $N_G^{r-r_u-1}(w)$, we know that distance of w from v in G is at most $r+r-r_u-1$ and thus every vertex in $N_G^{r-r_u-1}(w)$ is at distance at most $2r-r_u-1+r-r_u-1 < 3r$ from v in G. In other words, $N_{G'}^r(v) \subseteq N_G^{3r}(v)$ and thus $tw(N_{G'}^r(v)) \leq f(3r)+1$, as desired (again we add u to all bags). ☐

Corollary 1. *By adding a constant number of edges to a graph G having linear local treewidth, the resulting graph still has linear local treewidth.*

Corollary 2. *If a graph G has linear local treewidth, $tw(N_G^r(S))$ is linear in r for each set $S \subseteq V(G)$ of bounded size.*

Proof. We can add a new vertex v to G and attach it to all vertices in S. The result follows from Lemma 2 and the fact that $N_G^r(S) \subseteq N_G^{r+1}(v)$. ☐

The following simple lemma express the relation of local treewidth of a graph before and after contracting an edge.

Lemma 3. *Let G be a graph with $\mathrm{ltw}^{G_i}(r) \leq f(r)$ for all $r \in \mathbb{N}$. Let G' be a graph obtained by contracting an edge $\{u, w\}$. Then $\mathrm{ltw}^G(r) \leq f(r+1)$.*

Proof. The proof simply follows from the fact that for each vertex v, $N_{G'}^r(v) \subseteq N_G^{r+1}(v)$. ☐

Corollary 3. *By contracting a constant number of edges of a graph G which has linear local treewidth, the resulting graph still has linear local treewidth.*

Also we need to use the following theorem of Grohe [19] on local treewidth.

Theorem 9 ([19]). *Every apex-free h-almost embeddable graph G has linear local treewidth.*

Finally, we finish this section by stating the following lemma on treewidth after contracting all edges of a region of a bounded genus graph.

Lemma 4. *Contracting all edges of one of the faces of a bounded genus graph G embedded on a surface Σ changes its treewidth by at most three.*

Proof. According to a Theorem of Seymour and Thomas [39], the difference of treewidth of a bounded genus graph G and that of its dual is at most one. Contracting a face in G corresponds to deleting a vertex in its dual. Since deleting a vertex in the dual changes its treewidth by at most one, we obtain the desired result. ☐

We note that in the proof of Lemma 4, one could also use theorems of Robertson and Seymour (Theorem 6.6 of [25] and Theorem 2) which says the branchwidth of a graph and its dual are equal, and then use the the relation $\mathbf{bw}(G) \leq \mathrm{tw}(G)+1 \leq 3/2bw(G)$ [29] to show that the treewidth of a graph after the aforementioned contraction is at most $3/2tw(G)+3$. This is a weaker bound, but uses Theorems which are already published and it suffices for our goal.

7 Proof of Theorem 6

In this section, we show that any piece in the clique-sum decomposition of an apex-minor-free graph has linear local treewidth. Before starting the proof, it is worth mentioning that, in each piece G of the clique-sum decomposition of an H-minor-free graph $\hat{G}$, each vertex u of G has a corresponding vertex $\hat{u}$ in $\hat{G}$.[1] We say $\hat{u}$ is the *image* of u. However each edge $e = \{u, v\}$ in G might have a corresponding edge $\hat{e} = \{\hat{u}, \hat{v}\}$ in $\hat{G}$, in which case we call e an *actual edge* and call $\hat{e}$ the *image* of e, or might not have a corresponding edge in $\hat{G}$, in which case we call e a *virtual edge*. In the second case, the edge e is added via a clique-sum decomposition when completing the join set.

We start by "dealing with vortices" in such a way that we can omit them in the rest of the proof:

[1] For ease of notation, this section uses G to denote a piece of the graph instead of the original graph, denoted by $\hat{G}$ in this section.

Lemma 5. *Suppose G is a h-almost embeddable graph in a clique-sum decomposition of a graph $\hat{G}$. Let G' be the graph obtained by contracting all regions containing a vortex. Then graph G has linear local treewidth if and only if G' has linear local treewidth.*

Proof. As mentioned before, all edges of a face containing a vortex of G are actual edges. Now suppose $G = G_0 \cup G_1 \cup \cdots \cup G_h \cup X$ is almost-embedded in a surface Σ of genus g. Suppose U_i (as defined in Definition 2) is $\{u_i[1], u_i[2], \ldots, u_i[m_i]\}$. Let G'_0 be the graph obtained from G_0 by adding new vertices $c_1, c_2, \ldots, c_h$, placing c_i inside the region of G_0 in which G_i is placed, and adding edges $\{c_i, u_i[j]\}$ and $\{u_i[j], u_i[j+1]\}$ (where $j+1$ is treated modulo m_i) for all $1 \le i \le h$ and $1 \le j \le m_i$. By adding these edges, the vertices U_i, $1 \le i \le h$, form a cycle in G'_0. In addition, for each apex $x \in X$, we add an edge $\{x, c_i\}$ to G'_0 if and only if there is an edge $\{x, v\}$ in G for some $v \in G_i$, for all i. Using an argument similar to Grohe's argument [19] (in the proof of Theorem 9), we can show that if G'_0 has linear local treewidth $ar + b$, then G also has local treewidth at most $(h^2+1)ar + (h^2+1)b + O(h)$, which is linear in terms of r. (One difference is that, in Grohe's proof, there are no apices; however, by including all apices in all bags, we increase the treewidth by at most $O(h)$, and proximity is preserved because of the edges $\{x, c_i\}$.)

Next we construct graph G''_0 from G'_0 by subdividing each $u_i[j]$. More precisely, we obtain graph G''_0 from G_0 by adding new vertices c_i for $1 \le i \le h$ (again placed in the region of G_0 corresponding to G_i), vertices $u_i[j]$ and $u'_i[j]$ for $1 \le i \le h$ and $1 \le j \le m_i$, and edges $\{c_i, u'_i[j]\}$, $\{u'_i[j], u'_i[j+1]\}$, and $\{u_i[j], u'_i[j]\}$ (where $j+1$ is treated modulo m_i) for all $1 \le i \le h$ and $1 \le j \le m_i$. In addition, for each apex $x \in X$, we add an edge $\{x, c_i\}$ to G''_0 if and only if there is an edge $\{x, v\}$ in G for some $v \in G_i$. We can see that G'_0 has linear local treewidth if and only if G''_0 has linear local treewidth. More precisely, $\frac{1}{2}\mathrm{ltw}^{G''_0}(r) \le \mathrm{ltw}^{G'_0}(r) \le \mathrm{ltw}^{G''_0}(2r)$. Here the factor of 2 in the right-hand side appears because each edge in G'_0 is at most split in two in G''_0. The factor of $\frac{1}{2}$ in the left-hand side appears because the treewidth of each neighborhood in G''_0 grows by at most a factor of 2 with respect to the treewidth of the corresponding neighborhood in G'_0 (because we can replace each vertex $u_i[j]$ in a bag of a tree decomposition in G'_0 by both vertices $u_i[j]$ and $u'_i[j]$ in the corresponding bag of the tree decomposition in G''_0).

Finally, we obtain graph G'''_0 from G''_0 with the following transformation. First, we delete any edges $\{c_i, x\}$ for each $x \in X$. By Corollary 1, this operation preserves linear local treewidth, because there are at most h^2 such edges. Then we delete c_i and contract the face formed by vertices that used to be adjacent to c_i into a new vertex c'_i, for all $1 \le i \le h$. By Lemma 4 and because we preserve proximity, this operation preserves linear local treewidth. Finally, we delete c'_i and contract the face formed by vertices that used to be adjacent to c_i, which is the same as the original face of G_0 in which G_i was placed. For the same reason, this operation preserves linear local treewidth and proximity. The resulting graph G'''_0 is the G' desired by the lemma. □

Because we would like the piece G to be apex-minor-free like $\hat{G}$, we need to be able to obtain virtual edges via contractions. First we mention some basic (but usually hidden) information about virtual edges and where they arise in Theorem 1. The proof of Theorem 1 by Robertson and Seymour [38] shows that, in fact, the join set of each clique-sum contains at most three vertices of the piece that are neither apices nor vortices. In addition, by "taming the vortices" according to [34], we can strengthen Theorem 1 to say that all edges of a face containing a vortex are actual edges.

In the next step we deal with virtual edges. Intuitively, for each piece G in the clique-sum decomposition of the original graph $\hat{G}$, we construct a graph G^* which is a minor of $\hat{G}$ and "approximately" preserves the virtual edges of each piece in the following sense:

Definition 3. *Let G be a h-almost-embeddable graph in a clique-sum decomposition of a graph $\hat{G}$. The* approximation graph *of G, denoted by $\tilde{G}$, can be obtained as follows. First we perform the reduction described in Lemma 5, i.e., contract each vortex and its corresponding face into a single vertex. Next we remove the virtual edges from the graph and replace some of them as follows. For each clique-sum of G with a graph G' via a join set W, where $|W \cap (V(G) - X)| > 1$, we do the following:*

1. *If $|W \cap (V(G) - X)| = 2$, we add edges from all vertices of W to an arbitrary vertex in $W \cap (V(G) - X)$ (located on the surface).*
2. *If $|W \cap (V(G) - X)| = 3$ and there is more than one clique-sum that contains $W \cap (V(G) - X)$ in its join set, we add all edges among $W \cap (V(G) - X)$ and then add edges from all vertices in $W - (V(G) - X)$ to a vertex in $W \cap (V(G) - X)$. For each different clique-sum, we choose a different vertex in $W \cap (V(G) - X)$, as long as we can avoid repetition.*
3. *If $|W \cap (V(G) - X)| = 3$ and there is only one clique-sum that contains $W \cap (V(G) - X)$ in its join set, we add a new vertex v inside the triangle of $W \cap (V(G) - X)$ on the surface and then add all edges from all vertices of W to v.*

Lemma 6. *Let G be a h-almost-embeddable graph in a clique-sum decomposition of a graph $\hat{G}$. The approximation graph $\tilde{G}$ is a minor of $\hat{G}$.*

Proof. First, we note that after deleting vortices and contracting corresponding faces, the graph remains a minor of G. (As mentioned before, all edges of a face containing a vortex are actual edges.) By rewriting the initial clique-sum, we can ensure that the two terms in each clique-sum operation are connected even after removal of the join set; otherwise, we can rewrite into multiple clique-sum operations. Thus, if we have a clique-sum of G and G', G' can be contracted into a vertex v' adjacent to all vertices of the join set. Finally we can obtain the edges and vertices described in the definition of the approximation graph $\tilde{G}$ by at most one contraction per clique-sum operation. □

Lemma 7. *Let G be an h-almost-embeddable graph in a clique-sum decomposition of a graph $\hat{G}$. The approximation graph $\tilde{G}$ has linear local treewidth if and only if G has linear local treewidth.*

Proof. First, by Lemma 5, we only need to show that the graph G' in Lemma 5 has linear local treewidth if and only if the approximation graph $\tilde{G}$ has linear local treewidth. Note that $\tilde{G}$ is a subgraph of G' so, by a theorem of Hajiaghayi et al. [21, 13], $\mathrm{ltw}^{\tilde{G}}(r) \leq \mathrm{ltw}^{G'}(r)$. Finally, we can observe that $\mathrm{ltw}^{G'}(r) \leq 3\mathrm{ltw}^{\tilde{G}}(2r) + h$. The factor 2 appears because the distances in $\tilde{G}$ are at most twice the distances in G' (the construction preserves proximity). On the other hand, because we can afford to put all apices in all bags, no problem arises for edges between apices and other vertices. If for a clique-sum we have a triangle among vertices on the surface, we do not have problem either. The only case that can be problematic is when we have added one vertex v inside the triangle among vertices of W' on the surface. However, in this case, we can replace v by W' in all bags which contain v. Because of this construction, we obtain an additional factor of 3 in treewidth. One can easily observe that, after this construction, we have a tree-decomposition for G' as desired. □

To encourage reader, we point out that we are making progress. By construction, the approximation graph no longer contains vortices. Furthermore, by Lemmas 6 and 7, the approximation graph has linear local treewidth if (and only if) the piece G in the clique-sum decomposition of $\hat{G}$ has linear local treewidth, and simultaneously the approximation graph is a minor of the original graph $\hat{G}$. In the rest of the proof, we show that each apex $\tilde{G}$ can attach to only a bounded number of "local areas of planarity" which can be attached to the graph via clique-sums and themselves have linear local treewidth. To find such local areas, we first need to make sure that the representativity of the bounded-genus graph $\tilde{G} - X$ is relatively high, or it is embedded on a sphere. In the rest of the proof, because the approximation graph $\tilde{G}$ approximates the properties we need from the piece G, we use G and $\tilde{G}$ interchangeably.

Lemma 8. *Let G be a graph and $X \subseteq V(G)$ is a set of its apices where $|X| \leq h$, such that $G - X$ is 2-cell embedded in a surface Σ of genus g. We have one of the following cases:*

1. *Σ is a sphere, i.e. $g = 0$;*
2. *Representativity of G is at least N_1 (a positive constant to be determined later);*
3. *G has a set X' of size at most $h + N_1$ such that $G - X'$ can be 2-cell embedded in a surface with genus strictly smaller than g; and*

4. $G = G_1 \oplus G_2 \oplus \cdots \oplus G_k$ such that $X \subseteq V(G_i)$ for all i, $1 \le i \le k$, and $G_i - X$ can be 2-cell embedded in a surface with genus strictly smaller than g.

Proof. Suppose that the first and second cases do not apply. The proof follows from the definition of the representativity; if the representativity of G is less than N_1, then there exists a non-contractable noose C of size less than N_1. We remove these vertices from the surface, add them to X, and finally cut the surface via this noose. After this operation, either $G - X$ is connected, in which case by the definition of the cut the genus of $G - X$ is less than g, or $G - X$ has more than one component. In the former case, the third case applies. Finally, in the latter case, if $G - X$ has components $G_1, G_2, \ldots, G_k$, we define graph G'_i to be $G_i \cup C \cup X$. Then we add all edges among vertices $C \cup X$, forming a clique in G'_i, and write $G = G'_1 \oplus G'_2 \oplus \cdots \oplus G'_k$. Hence the fourth case applies. $\square$

By applying Lemma 8 any bounded number of times to the graph $\tilde{G}$, we can obtain a clique-sum decomposition $\tilde{G} = \tilde{G}_1 \oplus \tilde{G}_2 \oplus \cdots \oplus \tilde{G}_k$ such that each graph $\tilde{G}_i$ has a bounded number of apices $\tilde{X}_i$ and each $\tilde{G}_i - \tilde{X}_i$ is either a planar graph or a graph with representativity at least some constant lower bound. In fact, the representativity lower bound that we will need in Lemma 11 is not a fixed quantity, and indeed is different for each term $\tilde{G}_i$. Specifically, the lower bound is $l\tilde{h}_i N_2$ which depends on the number of apices $\tilde{h}_i = |\tilde{X}_i|$ in each term $\tilde{G}_i$, the degree l of the apex in the apex graph H, and the constant N_2 (determined in the next lemma). To achieve this lower bound, we apply Lemma 8 to each term $\tilde{G}_i$ that does not have sufficiently high representativity, unless the term is already planar. Note that, as we recurse, we increase the values of $\tilde{h}_i$ for some i and therefore require a stronger lower bound on the representativity of those terms. Because $\tilde{h}_i$ increases by only a constant, and because the maximum depth of recursion is bounded by g, the final values of $\tilde{h}_i$ are bounded, so the lower bound on representativity is also bounded.

Using Theorem 5, we only need to show that each resulting term $\tilde{G}_i$, $1 \le i \le k$, has linear local treewidth. We divide this claim into two cases: (1) where $\tilde{G}_i - \tilde{X}_i$ has sufficiently high representativity, and (2) where $\tilde{G}_i - \tilde{X}_i$ is planar. The next few lemmas deal with the first case.

First, we determine the necessary value of N_2:

Lemma 9. *For any apex graph H, there is an integer $N_2 \ge 0$ with the following property. Let G' be a minor of G such that $G' - X$ can be 2-cell embedded in a surface Σ of genus g, and let $\mathcal{T}$ be a tangle of $G' - X$ with metric d. Also assume that the representativity of $G' - X$ is at least N_2. Then, for each vertex $x \in X$, if x is adjacent to $v_1, \ldots, v_l$ where l is the degree of the apex in H, and $d(v_i, v_j) \ge N_2$ for all $1 \le i \ne j \le l$, then G has a minor isomorphic to H.*

Proof. The proof is very similar to the proof of [27, Section 9], [32, Theorems 4.4 and 4.5], or [35, Theorem 2.2], and we omit the details. $\square$

Using Lemma 9, we go one step towards finding the local areas of planarity as follows. Roughly, the following lemma says that the neighborhood of an apex in $\tilde{X}_i$ in the bounded-genus graph $\tilde{G}_i - \tilde{X}_i$ can be covered by a bounded number of bounded-radius disks, where the disks are defined by the metric d of the tangle $\tilde{G}_i - \tilde{X}_i$.

Lemma 10. *Let l be the degree of the apex in apex graph H which is excluded by the original graph $\hat{G}$, and let d be the metric of the tangle of $\tilde{G}_i - \tilde{X}_i$ mentioned in Lemma 9. Then, for each vertex x in the set $\tilde{X}_i$ of apices of $\tilde{G}_i$, there is a set C_i of at most l* centers *(vertices of $V(\tilde{G}_i) - \tilde{X}_i$) such that, for each neighbor u of x in $\tilde{G}_i - \tilde{X}_i$, there is a center c in C_i such that $d(c, u) \le N_2$.*

Proof. Comparing the terms $\tilde{G}_i$ to the approximation graph $\tilde{G}$, we have added edges only among apices $\tilde{X}_i$. If we ignore these edges among apices, $\tilde{G}_i$ becomes a minor of the approximation graph $\tilde{G}$ and thus a minor of the original graph $\hat{G}$, so it is H-minor-free. Now, the proof follows from applying Lemma 9 with $x \in \tilde{X}$ and $\tilde{G}_i - \tilde{X}_i$ as the minor. More precisely, we can greedily build the disk cover by repeatedly adding a disk of radius N_2 centered at a neighbor of x that is not already covered by the disks so far. When the cover is complete, the centers of the disks form a set C_i such that every pair of centers has distance at least N_2 (by construction). Therefore, by Lemma 9, C_i has size at most l. $\square$

We now combine this disk cover over all apices, and make the cover disjoint, to obtain our desired local areas of planarity as follows.

Lemma 11. *Let l be the degree of the apex in apex graph H which is excluded by the original graph $\hat{G}$, let d be the metric of the tangle of $\tilde{G}_i - \tilde{X}_i$ mentioned in Lemma 9, and let $\tilde{h}_i = |\tilde{X}_i|$. Suppose that the representativity of $\tilde{G}_i - \tilde{X}_i$ is at least $l\tilde{h}_i N_2$. Then there is a set C_i of at most $l\tilde{h}_i$ vertices such that, for each neighbor u of any $x \in \tilde{X}_i$ in $\tilde{G}_i - \tilde{X}_i$, there is exactly one center c in C_i for which $d(c, u) \leq l\tilde{h}_i N_2$.*

Proof. By Lemma 10, for each $x \in \tilde{X}_i$, we find the corresponding set of centers, and we let C_i be the union of all of them. Now, the only problem is that we might have some neighbors in $\tilde{G}_i - \tilde{X}_i$ of apices in $\tilde{X}_i$ that are near to (within radius of) more than one center in C_i (double coverage). Suppose one disk of radius r and center c intersects another disk of radius r' and center c'. We replace these disks by a single disk of radius $r + r'$ and centered at a vertex of distance r' from c and distance r from c'. (Such a vertex exists by the definition of the distance metric d.) Repeating this process, we eventually remove all intersections among disks. The maximum radius of any disk increases from the original maximum N_2 by at most a factor of the number of disks in the original C_i, which is at most $l\tilde{h}_i$. ☐

Now, we consider the structure of each of bounded-radius disks.

Lemma 12. *Let C be a bounded radius disk in $\tilde{G}_i$. More precisely, C contains all vertices and regions of distance at most r from the center $c \in V(G)$, where the distance is with respect to the metric of the tangle of $\tilde{G}_i - \tilde{X}_i$ mentioned in Lemma 9. Then we can write C as $(\leq N_1)$-sums of a graph of pathwidth at most N_2 and several planar graphs each with a bounded number of apices.*

Proof. The proof is very similar to the proofs in Sections 12 and 13 of Robertson and Seymour [26] and hence omitted in this paper. Intuitively, suppose we cut along the vertex-region paths from the center c to the vertices on the circumference of C (i.e., vertices of distance exactly r from c). Then the remaining pieces of C are planar, and the union of these cuts forms a bounded pathwidth structure. Furthermore, each of the vertex-region paths has bounded length, by the definition of the metric d, and so the shared boundary between a planar piece and the cut structure has bounded size. Therefore, we can recombine these pieces to form the disk C by bounded clique-sums. ☐

Lemma 12 also applies if we consider the apices of $\tilde{G}_i$, because we can easily add all apices to all planar graphs. If we also apply Lemma 12 once for each disk in the disjoint disk covering, then we arrive at the following:

Corollary 4. *Each graph $\tilde{G}_i$ can be written as $P_0 \oplus P_1 \oplus \cdots P_k$ where P_0 is a bounded-genus graph with a bounded number of vortices and apices such that each apex is attached only to vortices, and each P_i, $1 \leq i \leq k$, is a planar graph with a bounded number of apices. In addition, if we remove all edges between pairs of apices, the rest of the graph is a minor of $\tilde{G}_i$ and $\hat{G}$.*

At first glance, it may seem that we have arrived to our initial state before Lemma 5: again we have a bounded-genus graph plus apices and vortices. However, we have made substantial progress because now in this almost-embeddable graph P_0 all apices are attached only to vortices. Using a proof similar to the proof of Lemma 5, now we are able to remove apices and vortices from P_0 without destroying the linear local treewidth property. As a result, we convert P_0 into a bounded-genus graph without apices and vortices, and Eppstein [17] proved that such graphs have linear local treewidth.

To finish the high-representativity case, it remains only to show that each P_i, $1 \leq i \leq k$, has linear local treewidth. In fact, each P_i is a planar graph plus apices, making them equivalent to the other case of $\tilde{G}_i$'s in which $\tilde{G}_i - \tilde{X}_i$ is planar. Therefore, we can finish all remaining loose ends with a final lemma:

Lemma 13. *If a graph G has a subset X of vertices of size at most a constant h such that (1) if we remove edges between pairs of vertices in X, G becomes H-minor-free where H is an apex graph, and (2) $G - X$ is a planar graph, then G has linear local treewidth.*

Proof. First, we can assume that all edges among vertices in X exist in graph G, because if we add all such edges, the new graph has linear local treewidth if and only if the original graph has linear local treewidth by Corollary 1 and because $|X| \le h$.

We claim that it suffices to show that $\mathrm{tw}(G[N_r^G(X)]) \le cr + d$ for some constants c and d. Suppose we had this relation. Because $N_r^G(v) \subseteq N_r^G(X)$ for each $v \in X$, we have $\mathrm{tw}(G[N_r^G(v)]) \le \mathrm{tw}(G[N_r^G(X)]) \le cr + d$ for each $v \in X$. For each $v \in V(G) - X$, if $N_r^G(v) \cap X = \emptyset$, then because G is a planar graph, we have $\mathrm{tw}(G[N_r^G(v)]) \le 3r - 1$ which is linear. On the other hand, if $v \in V(G) - X$ and there exists an $x \in N_r^G(v) \cap X$, then we have $N_r^G(v) \subseteq N_{2r}^G(x)$ and thus $\mathrm{tw}(G[N_r^G(v)]) \le \mathrm{tw}(G[N_{2r}^G(X)]) \le 2cr + d$ as desired.

To show that $\mathrm{tw}(G[N_r^G(X)]) \le cr + d$, we consider a fixed r and let $w = \mathrm{tw}(G[N_r^G(X)])$. We know that $w - h \le \mathrm{tw}(G[N_r^G(X)] - X) \le w$, so $\mathrm{tw}(G[N_r^G(X)] - X) = \Theta(w)$. Because the planar graph $G[N_r^G(X)] - X$ has treewidth at least $\Theta(w)$, by [36, Theorem 6.2], it has a $(\Theta(w) \times \Theta(w))$-grid as a minor. If, during the course of taking a minor to obtain such a grid, we apply only contractions and ignore deletions, then we obtain a partially triangulated $(\Theta(w) \times \Theta(w))$ grid R.

Next we consider the tangle of the grid R and its corresponding metric. Let x be a vertex in X and let $N(x)$ be the neighbors of x in R (i.e., we consider the neighbors of x after the contractions). By the theory of Robertson and Seymour (see [27, Section 9] or [32, Theorems 4.4 and 4.5]), if we have n vertices of an $(r \times r)$-grid such that each pair has distance at least some constant N_2 and each vertex has distance at least N_2 from the boundary of the outer region, then by taking a minor of the $(r \times r)$-grid we can construct any planar graph whose vertex set is precisely these n vertices. (In fact, this theorem uses the distance metric of the tangle, but for triangulated grids, this distance is proportional to the normal graph distance.) Because G without edges among X is H-minor-free and H is an apex graph, n cannot be more than $|V(H)| - 1$. In other words, for each vertex $x \in X$, there is a set $C_x \subseteq N(x)$ of at most $h_1 = |V(H)| - 1$ vertices (centers) such that the distance in R between any pair of vertices in C_x is at least N_2, and every vertex of $N(x)$ in the central $(\Theta(w) - 2N_2 \times \Theta(w) - 2N_2)$-grid R' of R has distance at most N_2 from one of the vertices of C_x.

Consider the set of disks with centers from C_x over all apices $x \in X$, all with radius N_2. By "merging" overlapping disks as in Lemma 11, we obtain a bounded number of disjoint disks with bounded radius.

We perform the following modifications to the planar graph $G - X$. For every two adjacent rows that intersect a disk, we contract vertically into a single row, and then remove edges to avoid multi-edges and loops. (These edge removals are valid because they do not affect distances.) Similarly, we contract adjacent columns that intersect disks. We also contract the N_2 outermost rings of the grid down to a single ring (again removing edges to avoid multi-edges and loops). Finally, we add edges to connect one vertex v on the new outermost ring to every other vertex on the outermost ring, and add v to the set of centers. The resulting graph R'' is still planar; indeed, it is a partially triangulated $(\Theta(w) \times \Theta(w))$-grid, and thus its treewidth is still $\Theta(w)$. Furthermore, each disk has been contracted down to a single vertex, which we can think of as the center.

Because every vertex in the original grid R has distance at most r from an apex, every vertex in the new grid R'' has distance at most r from a center. By Corollary 2, the treewidth is linear in r, and therefore $\Theta(w) = \Theta(r)$. □

Thus by Lemma 13 we conclude that each term $\tilde{G}_i$ and thus the approximation graph $\tilde{G}$ and G itself (by Lemma7) have linear local treewidth. Because the original graph $\hat{G}$ can be written as a clique-sum of such graphs of linear local treewidth, by Lemma 5, the apex-minor-free graph $\hat{G}$ has linear local treewidth.

8 Conclusions and Future Work

In this paper, by reproving the main theorem of Eppstein [17] using the deep Graph Minor Theory of Robertson and Seymour, we showed that the concept of local treewidth and linear local treewidth are the same for minor-closed families of graphs. Using this result, we obtain PTASs with approximation ratio $1+\varepsilon$ running in $2^{O(1/\varepsilon)}n^{O(1)}$ time (instead of $2^{2^{2^{O(1/\varepsilon)}}}n^{O(1)}$ time) for many NP-complete problems on minor-closed classes of graphs with bounded local treewidth.

The constants that we obtain for linear local treewidth of apex-minor-free graphs are not the best. It would be interesting to improve these constants, even for special class of graphs (see e.g. Demaine et al. [13], Grohe [19], or Eppstein [19]), because of the direct improvement on the exponents in the running time of (and thus the practicality of) PTASs for many NP-complete problems on these graphs.

We re-enforce another open problem posed by Eppstein [17] which asks whether there are natural nontrivial families of graphs which are not minor-closed but that have bounded local treewidth. (A trivial example is bounded-degree graphs, or other classes in which a bound on diameter imposes a limit on total graph size.) Given our results, a natural question is whether there are any (non-minor-closed) families of graphs that have bounded local treewidth but not linear local treewidth.

Recently, several papers have been published on exponential speedup of fixed parameter algorithms on special class of graphs; see e.g. [1, 12, 15]. Most of these results bound the treewidth of the graph as a linear function in the square-root of the size of a minimum dominating set. One can easily observe that there is no such bound for general apex graphs, and therefore the most general minor-closed families of graphs to which we could hope to obtain such a relation are apex-minor-free graphs. We obtain that indeed apex-minor-free graphs have such a relation, by following through our series of reductions and by showing that, at each step except the last, we preserve linear dominating set number in the same way that we preserve linear local treewidth, and finally in the planar case we obtain square-root bound. In fact, we conjecture that any apex-minor-free graph can be represented as a clique-sum of almost-embeddable graphs such that, in each piece, an apex is attached only to vortices. This conjectured characterization of apex-minor-free graphs would easily lead to our linear local treewidth bound and the aforementioned bound on treewidth with respect to dominating set.

Acknowledgments

The authors are indebted to Paul D. Seymour for many important discussions that led to the proof of the main combinatorial result of this paper (Theorem 6) and for providing a portal into the Graph Minor Theory and revealing some of its hidden structure that we use in this proof. We also thank Fedor Fomin, Naomi Nishimura, Prabhakar Ragde, and Dimitrios Thilikos for encouragement and helpful discussions.

References

1. Jochen Alber, Hans L. Bodlaender, Henning Fernau, and Rolf Niedermeier. Fixed parameter algorithms for planar dominating set and related problems. In *Algorithm theory—Scandinavian Workshop on Algorithm Theory 2000 (Bergen, 2000)*, pages 97–110. Springer, Berlin, 2000.
2. Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for graphs with excluded minor and its applications. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (Baltimore, MD, 1990)*, pages 293–299, 1990.
3. Eyal Amir. Efficient approximation for triangulation of minimum treewidth. In *Uncertainty in Artificial Intelligence: Proceedings of the Seventeenth Conference (UAI-2001)*, pages 7–15, San Francisco, CA, 2001. Morgan Kaufmann Publishers.
4. Stefan Arnborg and Andrzej Proskurowski. Linear time algorithms for NP-hard problems restricted to partial k-trees. *Discrete Appl. Math.*, 23(1):11–24, 1989.
5. Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. Assoc. Comput. Mach.*, 41(1):153–180, 1994.
6. Hans L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996.
7. John A. Bondy and U. S. R. Murty. *Graph Theory with Applications.* American Elsevier Publishing Co., Inc., New York, 1976.

8. Moses Charikar and Amit Sahai. Dimension reduction in the l_1 norm. In *The 43th Annual Symposium on Foundations of Computer Science (FOCS'02)*, pages 551–560, 2002.
9. Chandra Chekuri, Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Sinclair Alistair. Embedding k-outerplanar graphs into ℓ_1. In *ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*, pages 527 – 536, 2003.
10. Zhi-Zhong Chen. Efficient approximation schemes for maximization problems on $K_{3,3}$-free or K_5-free graphs. *J. Algorithms*, 26(1):166–187, 1998.
11. Norishige Chiba, Takao Nishizeki, and Nobuji Saito. An approximation algorithm for the maximum independent set problem on planar graphs. *SIAM J. Comput.*, 11(4):663–675, 1982.
12. Erik D. Demaine, Fedor V. Fomin, Mohammadtaghi Hajiaghayi, and Dimitrios M. Thilikos. Fixed-parameter algorithms for the (k, r)-center in planar graphs and map graphs. In *The 30th International Colloquium on Automata, Languages and Programming (ICALP 2003)*. 2003. To appear.
13. Erik D. Demaine, Mohammadtaghi Hajiaghayi, Naomi Nishimura, Prabhakar Ragde, and Dimitrios M. Thilikos. Approximation algorithms for classes of graphs excluding single-crossing graphs as minors. Submitted to Journal of Computer and System Sciences.
14. Erik D. Demaine, Mohammadtaghi Hajiaghayi, and Dimitrios M. Thilikos. A 1.5-approximation for treewidth of graphs excluding a graph with one crossing. In *The 5th International Workshop on Approximation Algorithms for Combinatorial Optimization (Italy, APPROX 2002), LNCS*, pages 67–80. 2002.
15. Erik D. Demaine, Mohammadtaghi Hajiaghayi, and Dimitrios M. Thilikos. Exponential speedup of fixed parameter algorithms on $K_{3,3}$-minor-free or K_5-minor-free graphs. In *The 13th Annual International Symposium on Algorithms and Computation (Vancouver, ISAAC 2002)*, pages 262–273. 2002.
16. David Eppstein. Subgraph isomorphism in planar graphs and related problems. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, CA, 1995)*, pages 632–640, New York, 1995. ACM.
17. David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3-4):275–291, 2000.
18. Markus Frick and Martin Grohe. Deciding first-order properties of locally tree-decomposable graphs. *J. ACM*, 48(6):1184–1206, 2001.
19. Martin Grohe. Local tree-width, excluded minors, and approximation algorithms. To appear in Combinatorica.
20. Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Sinclair Alistair. Cuts, trees and l_1-embeddings of graphs. In *The 40th Annual Symposium on Foundations of Computer Science (FOCS'99)*, pages 399–409, 1999.
21. Mohammadtaghi Hajiaghayi, Naomi Nishimura, Prabhakar Ragde, and Dimitrios M. Thilikos. Fast approximation schemes for $K_{3,3}$-minor-free or K_5-minor-free graphs. In *Euroconference on Combinatorics, Graph Theory and Applications 2001 (Barcelona, 2001)*. 2001.
22. Philip N. Klein, Serge A. Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *The Twenty-Fifth Annual ACM Symposium on Theory of Computing (STOC'93)*, pages 682–690, 1993.
23. Richard J. Lipton and Robert Endre Tarjan. Applications of a planar separator theorem. *SIAM J. Comput.*, 9(3):615–627, 1980.
24. Serge A. Plotkin, Satish Rao, and Warren D. Smith. Shallow excluded minors and improved graph decompositions. In *The Fifth Annual ACM-SIAM Symposium on Discrete Algorithms(SODA'94)*, pages 462–470, 1994.
25. Neil Robertson. Graph minors. XI. Circuits on a surface. *J. Combin. Theory Ser. B*, 60(1):72–106, 1994.
26. Neil Robertson and P. D. Seymour. Graph minors. XIX. Well-quasi-ordering on a surface. To appear in J. Combin. Theory Ser. B.
27. Neil Robertson and P. D. Seymour. Graph minors. VII. Disjoint paths on a surface. *J. Combin. Theory Ser. B*, 45(2):212–254, 1988.
28. Neil Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991.
29. Neil Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B*, 52(2):153–190, 1991.
30. Neil Robertson and P. D. Seymour. Graph minors. XI. Circuits on a surface. *J. Combin. Theory Ser. B*, 60(1):72–106, 1994.

31. Neil Robertson and P. D. Seymour. Graph minors. XII. Distance on a surface. *J. Combin. Theory Ser. B*, 64(2):240–272, 1995.
32. Neil Robertson and P. D. Seymour. Graph minors. XII. Distance on a surface. *J. Combin. Theory Ser. B*, 64(2):240–272, 1995.
33. Neil Robertson and P. D. Seymour. Graph minors. XVII. Taming a vortex. *J. Combin. Theory Ser. B*, 77(1):162–210, 1999.
34. Neil Robertson and P. D. Seymour. Graph minors. XVII. Taming a vortex. *J. Combin. Theory Ser. B*, 77(1):162–210, 1999.
35. Neil Robertson and Paul Seymour. Excluding a graph with one crossing. In *Graph structure theory (Seattle, WA, 1991)*, pages 669–675. Amer. Math. Soc., Providence, RI, 1993.
36. Neil Robertson, Paul Seymour, and Robin Thomas. Quickly excluding a planar graph. *J. Combin. Theory Ser. B*, 62(2):323–348, 1994.
37. Neil Robertson and Paul D. Seymour. Graph minors. II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
38. Neil Robertson and Paul D. Seymour. Graph minors XVI. excluding a non-planar graph. 2003. To appear in J. Combin. Theory Ser. B, available electronically in journal site.
39. P.D. Seymour and Robin Thomas, 2003. Unpublished manuscript.
40. Mihalis Yannakakis. Node- and edge-deletion NP-complete problems. In *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, CA, 1978)*, pages 253–264. ACM press, New York, 1978.