

Bidimensionality and Geometric Graphs

Fedor V. Fomin* Daniel Lokshtanov† Saket Saurabh‡

Abstract

Bidimensionality theory was introduced by Demaine et al. [*JACM 2005*] as a framework to obtain algorithmic results for hard problems on minor closed graph classes. The theory has been successfully applied to yield subexponential time parameterized algorithms, EPTASs and linear kernels for many problems on families of graphs excluding a fixed graph H as a minor. In this paper we use several of the key ideas from Bidimensionality to give a new generic approach to design EPTASs and subexponential time parameterized algorithms for problems on classes of graphs which are not minor closed, but instead exhibit a geometric structure. In particular we present EPTASs and subexponential time parameterized algorithms for FEEDBACK VERTEX SET, VERTEX COVER, CONNECTED VERTEX COVER, DIAMOND HITTING SET, on map graphs and unit disk graphs, and for CYCLE PACKING and MINIMUM-VERTEX FEEDBACK EDGE SET on unit disk graphs. To the best of our knowledge, these results were previously unknown, with the exception of the EPTAS and a subexponential time parameterized algorithm on unit disk graphs for VERTEX COVER, which were obtained by Marx [*ESA 2005*] and Alber and Fiala [*J. Algorithms 2004*], respectively.

Our results are based on the recent decomposition theorems proved by Fomin et al. in [*SODA 2011*] and novel grid-excluding theorems in unit disc and map graphs without large cliques. Our algorithms work directly on the input graph and do not require the geometric representations of the input graph. We also show that our approach can not be extended in its full generality to more general classes of geometric graphs, such as intersection graphs of unit balls in \mathbb{R}^d , $d \geq 3$. Specifically, we prove that FEEDBACK VERTEX SET on unit-ball graphs in \mathbb{R}^3 neither admits PTASs unless $P=NP$, nor subexponential time algorithms unless the Exponential Time Hypothesis fails. Additionally, we show that the decomposition theorems which our approach is based on fail for disk graphs and that therefore any extension of our results to disk graphs would require new algorithmic ideas. On the other hand, we prove that our EPTASs and subexponential time algorithms for VERTEX COVER and CONNECTED VERTEX COVER carry over both to disk graphs and to unit-ball graphs in \mathbb{R}^d for every fixed d .

*Department of Informatics, University of Bergen, Norway.

fedor.fomin@ii.uib.no

†Department of Computer Science and Engineering, University of California, San Diego, USA.

dlokshtanov@cs.ucsd.edu

‡The Institute of Mathematical Sciences, Chennai, India.

saket@imsc.res.in

1 Introduction

Algorithms for hard optimization problems on intersection graphs of systems of geometric objects is a well studied area in Computer Science motivated by various applications in wireless networks [32], computational biology [42] and map labeling [1]. While most problems remain NP-complete even when restricted to such classes, the restriction of a problem to a geometric class is usually much more tractable algorithmically than the unrestricted problem. For example, for planar graphs and more generally minor closed families of graphs, the Bidimensionality theory of Demaine et al. [17] simultaneously demonstrates the tractability of most natural problems with respect to subexponential time parameterized algorithms [17], approximation [18, 26] and kernelization [27]. For related “geometric” classes of graphs that are not closed under taking minors, the picture is considerably more heterogenous and the situation is less understood. The objective of this article is to take a step towards clearing the picture for geometric graph classes.

Most of the known approximation schemes that have been obtained for graph problems on geometric graph classes use a variation of the well-known *shifting* technique introduced in the classical works of Baker [5] and of Hochbaum and Maass [28]. Hunt et al. [29] used the shifting technique to give polynomial time approximation schemes (PTASs) for a number of problems such as MAXIMUM INDEPENDENT SET and MINIMUM DOMINATING SET on unit disk graphs and λ -precision disk graphs. Independently, Erlebach et al. [24] and Chan [9] generalized the shifting technique and gave PTASs for MAXIMUM INDEPENDENT SET and MINIMUM VERTEX COVER on disk graphs and on intersection graphs of fat objects. Marx in [34] obtained an efficient polynomial time approximation schemes (EPTAS) for MINIMUM VERTEX COVER on unit disk graphs. Chen in [10] and Demaine et al. [16] used similar approaches to obtain a PTAS for MAXIMUM INDEPENDENT SET and MINIMUM r -DOMINATING SET on map graphs. One of the known limitations of the shifting technique is that it generally only applies to local problems such as VERTEX COVER and variants of DOMINATING SET, and fails for non-local problems such as FEEDBACK VERTEX SET and CYCLE PACKING. For problems on planar and H -minor-free graphs, Bidimensionality is able to handle non-locality by applying *treewidth* based decomposition. It is tempting to ask whether treewidth based decomposition can be useful for other graph classes as well.

In this article we use key ideas from Bidimensionality and design a general approach that can be used to give EPTASs and subexponential time parameterized algorithms for many problems on map graphs and for unit disk graphs, and in some cases on even more general geometric classes of graphs. We present EPTASs and subexponential time parameterized algorithms for FEEDBACK VERTEX SET, VERTEX COVER, CONNECTED VERTEX COVER, DIAMOND HITTING SET, MINIMUM-VERTEX FEEDBACK EDGE SET on map and unit disk graphs, and for CYCLE PACKING and MINIMUM-VERTEX FEEDBACK EDGE SET on unit disk graphs. Our approach is based on the concept of *truly sublinear treewidth*, recently introduced by the authors in [26] as a tool to give EPTASs for bidimensional problems on minor closed graph classes. Roughly speaking, a graph class has truly sublinear treewidth if adding k vertices to any graph in the class such that the resulting graph is in the class as well, increases its treewidth by $O(k^\epsilon)$ for $\epsilon < 1$. The techniques in [26] can not be applied directly to map graphs and unit disk graphs, because both graph classes contain arbitrarily large cliques and hence do not have truly sublinear treewidth. We overcome this obstacle by showing that cliques are the only pathological case. Namely, we prove that map graphs and unit disk graphs that exclude large clique subgraphs have truly sublinear treewidth. Our EPTASs work in two steps, first we “clean” the input graph for large cliques, and then we apply the decomposition theorems from [26].

The initial application of Bidimensionality was in the design of subexponential parameterized algorithms on planar, and more generally, on H -minor-free graphs [17]. Demaine et al. [16] used Bidimensionality to obtain subexponential parameterized algorithms for DOMINATING SET, and more generally, for (k, r) -CENTER on map graphs. We show that after “cleaning” unit disk and map graphs from large cliques, it is possible to use Bidimensionality to solve many parameterized problems in subexponen-

tial time on these classes of graphs. To the best of our knowledge, prior to our work parameterized subexponential algorithms on unit disk graphs were known only for VERTEX COVER [2]. The important ingredient of our algorithms are the analogues of excluding grid theorems of Robertson et al. for planar graphs [38] and of Demaine and Hajiaghayi for H -minor free graphs [19]. We show that the treewidth of every unit disc or map graph excluding a clique of constant size as a subgraph and excluding a $k \times k$ grid as a minor, is $\mathcal{O}(k)$.

Our algorithms do not require geometric representations of the input graphs. Since recognition of unit disk graphs is NP-hard [13] and the exponent of the polynomial bounding the running time of map graph recognition algorithm is about 120 [40], the robustness of our algorithms is a serious advantage.

We explore to which degree our approach can be lifted to other classes of graphs. Our investigations show that it is unlikely that the full power of our approach can be generalized to disk graphs or to unit ball graphs in \mathbb{R}^d — intersection graphs of unit-balls in \mathbb{R}^d , $d \geq 3$. Specifically we prove that FEEDBACK VERTEX SET on unit-ball graphs in \mathbb{R}^3 neither admits a PTASs unless $P=NP$, nor a subexponential time algorithm unless the Exponential Time Hypothesis fails. Furthermore we show that disk graphs which exclude the clique on four vertices as a subgraph do not have truly sublinear treewidth. On the other hand, an adaptation of our techniques yields EPTASs and subexponential time parameterized algorithms for VERTEX COVER and CONNECTED VERTEX COVER both on disk graphs and on unit disk graphs in \mathbb{R}^d for every fixed integer dimension d .

A natural question is whether our results can be extended to handle larger classes of problems on map graphs and unit disk graphs. It appears that the main obstacle to generalizing our approach is to design more general clique cleaning procedures. In particular, Marx [34] showed that DOMINATING SET and INDEPENDENT SET are $W[1]$ -hard even on unit disk graphs. This means that the two problems neither admit EPTASs nor FPT algorithms unless $FPT=W[1]$, a complexity collapse considered very unlikely. While no clique cleaning procedure is known for these two problems, a simple modification of our techniques show that both problems admit both EPTASs and subexponential time algorithms on map graphs and unit disk graphs excluding large cliques as subgraphs. Thus it seems that we are able to handle exactly the problems for which cliques can be removed efficiently.

2 Definitions and Notations

In this section we give various definitions which we make use of in the paper. Let G be a graph with vertex set $V(G)$ and edge set $E(G)$. A graph G' is a *subgraph* of G if $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$. The subgraph G' is called an *induced subgraph* of G if $E(G') = \{uv \in E(G) \mid u, v \in V(G')\}$, in this case, G' is also called the subgraph *induced by* $V(G')$ and denoted by $G[V(G')]$. For a vertex set S , by $G \setminus S$ we denote $G[V(G) \setminus S]$. A graph class \mathcal{G} is *hereditary* if for any graph $G \in \mathcal{G}$ all induced subgraphs of G are in \mathcal{G} . By $N(u)$ we denote (open) neighborhood of u , that is, the set of all vertices adjacent to u . Similarly, by $N[u] = N(u) \cup \{u\}$ we define the closed neighborhood. The degree of a vertex v in G is $|N_G(v)|$. We denote by $\Delta(G)$ the maximum vertex degree in G . For a subset $D \subseteq V(G)$, we define $N[D] = \cup_{v \in D} N[v]$ and $N(D) = N[D] \setminus D$. Given an edge $e = xy$ of a graph G , the graph G/e is obtained from G by contracting the edge e . That is, the endpoints x and y are replaced by a new vertex v_{xy} which is adjacent to the old neighbors of x and y (except from x and y). A graph H obtained by a sequence of edge-contractions is said to be a *contraction* of G . A graph H is a *minor* of a graph G if H is the contraction of some subgraph of G and we denote it by $H \leq_m G$. We also use the following equivalent characterization of minors.

Proposition 1 ([21]). *A graph H is a minor of G if and only if there is a map $\phi : V(H) \rightarrow 2^{V(G)}$ such that for every vertex $v \in V(H)$, $G[\phi(v)]$ is connected, for every pair of vertices $v, u \in V(H)$, $\phi(u) \cap \phi(v) = \emptyset$, and for every edge $uv \in E(H)$, there is an edge $u'v' \in E(G)$ such that $u' \in \phi(u)$ and $v' \in \phi(v)$.*

Let G, H be two graphs. A subgraph G' of G is said to be a *minor-model* of H in G if G' contains H as a minor. The $(r \times r)$ -grid is the Cartesian product of two paths of lengths $r - 1$.

Treewidth A *tree decomposition* of a graph G is a pair (\mathcal{X}, T) , where T is a tree and $\mathcal{X} = \{X_i \mid i \in V(T)\}$ is a collection of subsets of V such that the following conditions are satisfied.

1. $\bigcup_{i \in V(T)} X_i = V(G)$.
2. For each edge $xy \in E(G)$, $\{x, y\} \subseteq X_i$ for some $i \in V(T)$.
3. For each $x \in V(G)$ the set $\{i \mid x \in X_i\}$ induces a connected subtree of T .

Each X_i is called the bag of a tree decomposition. The *width* of the tree decomposition is $\max_{i \in V(T)} |X_i|$.
 1. The *treewidth* of a graph G , $\text{tw}(G)$, is the minimum width over all tree decompositions of G .

Plane, unit disk and map graphs. In this paper we use the expression *plane graph* for any planar graph drawn in the Euclidean plane \mathbb{R}^2 without any edge crossing. We do not distinguish between a vertex of a plane graph and the point of \mathbb{R}^2 used in the drawing to represent the vertex or between an edge and the curve representing it. We also consider plane graph G as the union of the points corresponding to its vertices and edges. We call by *face* of G any connected component of $\mathbb{R}^2 \setminus (E(G) \cup V(G))$. The *boundary* of a face is the set of edges incident to it. If the boundary of a face f forms a cycle then we call it a *cyclic face*. A *disk graph* is the intersection graph of a family of (closed) disks in \mathbb{R}^2 . A *unit disk graph* is the intersection graph of a family of unit disks in \mathbb{R}^2 . The notion of a map graph is due to Chen et al. [11]. A *map* \mathcal{M} is a pair (\mathcal{E}, ω) , where \mathcal{E} is a plane graph and each connected component of \mathcal{E} is biconnected, and ω is a function that maps each face f of \mathcal{E} to 0 or 1 in a way that whenever $\omega(f) = 1$, f is a closed face. A face f of \mathcal{E} is called *nation* if $\omega(f) = 1$, *lake* otherwise. The graph associated with \mathcal{M} is a simple graph G , where $V(G)$ consists of the nations on \mathcal{M} and $E(G)$ consists of all $f_1 f_2$ such that f_1 and f_2 are adjacent (that is shares at least one vertex). We call G a *map graph*. By $N(\mathcal{E})$ we denote the set of nations of \mathcal{E} .

Counting Monadic Second Order Logic. The syntax of MSO of graphs includes the logical connectives $\vee, \wedge, \neg, \Leftrightarrow, \Rightarrow$, variables for vertices, edges, set of vertices and set of edges, the quantifiers \forall, \exists that can be applied to these variables, and the following five binary relations:

1. $u \in U$ where u is a vertex variable and U is a vertex set variable.
2. $d \in D$ where d is an edge variable and D is an edge set variable.
3. $\text{inc}(d, u)$, where d is an edge variable, u is a vertex variable, and the interpretation is that the edge d is incident on the vertex u .
4. $\text{adj}(u, v)$, where u and v are vertex variables, and the interpretation is that u and v are adjacent.
5. Equality of variables, $=$, representing vertices, edges, set of vertices and set of edges.

Counting monadic second-order logic (CMSO) is monadic second-order logic (MSO) additionally equipped with an atomic formula $\text{card}_{n,p}(U)$ for testing whether the cardinality of a set U is congruent to n modulo p , where n and p are integers independent of the input graph such that $0 \leq n < p$ and $p \geq 2$. We refer to [4, 14, 15] for a detailed introduction to CMSO. MIN-CMSO and MAX-CMSO problems are graph optimization problems where the objective is to find a maximum or minimum sized vertex or edge set satisfying a CMSO-expressible property. In particular, in a MIN/MAX-CMSO graph problem Π we are given a graph G as input. The objective is to find a minimum/maximum cardinality vertex/edge set S such that the CMSO-expressible predicate $P_\Pi(G, S)$ is satisfied.

Bidimensionality and Separability. Our results concern graph optimization problems where the objective is to find a vertex or edge set that satisfies a feasibility constraint and maximizes or minimizes a problem-specific objective function. For a problem Π and vertex (edge) set S let $\phi_\Pi(G, S)$ be the feasibility constraint returning **true** if S is feasible and **false** otherwise. Let $\kappa_\Pi(G, S)$ be the objective function. In most cases, $\kappa_\Pi(G, S)$ will return $|S|$. We will only consider problems where every instance has at least one feasible solution. Let \mathcal{U} be the set of all graphs. For a graph optimization problem Π

let $\pi : \mathcal{U} \rightarrow \mathbb{N}$ be a function returning the objective function value of the optimal solution of Π on G . We say that a problem Π is *minor-closed* if $\pi(H) \leq \pi(G)$ whenever H is a minor of G . We now define bidimensional problems.

Definition 1 ([17]). *A graph optimization problem Π is minor-bidimensional if Π is minor-closed and there is $\delta > 0$ such that $\pi(R) \geq \delta r^2$ for the $(r \times r)$ -grid R . In other words, the value of the solution on R should be at least δr^2 .*

Demaine and Hajiaghayi [18] define the *separation* property for problems, and show how separability together with bidimensionality is useful to obtain EPTASs on H -minor-free graphs. In our setting a slightly weaker notion of separability is sufficient. In particular the following definition is a reformulation of the requirement 3 of the definition of separability in [18] and similar to the definition used in [27] to obtain kernels for bidimensional problems.

Definition 2. *A minor-bidimensional problem Π has the separation property if given any graph G and a partition of $V(G)$ into $L \uplus S \uplus R$ such that $N(L) \subseteq S$ and $N(R) \subseteq S$, and given an optimal solution OPT to G , $\pi(G[L]) \leq \kappa_{\Pi}(G[L], OPT \cap L) + \mathcal{O}(|S|)$ and $\pi(G[R]) \leq \kappa_{\Pi}(G[R], OPT \cap R) + \mathcal{O}(|S|)$.*

In Definition 2 we slightly misused notation. Specifically, in the case that OPT is an edge set we should not be considering $OPT \cap R$ and $OPT \cap L$ but $OPT \cap E(G[R])$ and $OPT \cap E(G[L])$ respectively.

Reducibility, η -Transversability and Graph Classes with Truly Sublinear Treewidth. We now define three of the central notions of this article.

Definition 3. *A graph optimization problem Π with objective function κ_{Π} is called reducible if there exist a MIN/MAX-CMSO problem Π' and a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that*

1. *there is a polynomial time algorithm that given G and $X \subseteq V(G)$ outputs G' such that $\pi'(G') = \pi(G) \pm \mathcal{O}(|X|)$ and $\mathbf{tw}(G') \leq f(\mathbf{tw}(G \setminus X))$,*
2. *there is a polynomial time algorithm that given G and $X \subseteq V(G)$, G' and a vertex (edge) set S' such that $P_{\Pi'}(G', S')$ holds, outputs S such that $\phi_{\Pi}(G, S) = \mathbf{true}$ and $\kappa_{\Pi}(G, S) = |S'| \pm \mathcal{O}(|X|)$.*

Definition 4. *A graph optimization problem Π is called η -transversable if there is a polynomial time algorithm that given a graph G outputs a set X of size $\mathcal{O}(\pi(G))$ such that $\mathbf{tw}(G \setminus X) \leq \eta$.*

Definition 5. *Graph class \mathcal{G} has truly sublinear treewidth with parameter λ , $0 < \lambda < 1$, if for every $\eta > 0$, there exists $\beta > 0$ such that for any graph $G \in \mathcal{G}$ and $X \subseteq V(G)$ the condition $\mathbf{tw}(G \setminus X) \leq \eta$ yields that $\mathbf{tw}(G) \leq \eta + \beta |X|^{\lambda}$.*

Throughout this paper K_t denotes a complete graph on t vertices and we say that a graph G is K_t -free if G does not contain K_t as an induced subgraph.

3 Structure of K_t -free Geometric Graphs

In this section we show that unit disk graphs and map graphs with bounded maximum clique size has truly sublinear treewidth.

3.1 Structure of K_t -free Unit Disk Graphs.

We start with the following well known observation (see [33, Lemma 3.2]) about K_t -free unit disk graphs.

Observation 1. *If a unit disk graph G is K_t -free then $\Delta(G) \leq 6t$.*




Figure 1: (A) Drawing of a planar (multi)graph P_I formed by the drawing of four disks; (B) The dual graph D_I of P_I ; (C) The graph P_G .

Observation 1 allows us to prove theorems on unit disk graphs of bounded maximum degree and then use these results for K_t -free graphs.

Let G be a unit disk graph generated by $\mathcal{B} = \{B_1, \dots, B_n\}$ and $\Delta(G) = \Delta$. We will associate an auxiliary planar graph P_G with G such that the treewidth of these two graphs is linearly related. Let P_I be a planar graph defined as follows. Consider the embedding (drawing) of the unit disks $\mathcal{B} = \{B_1, \dots, B_n\}$ in the plane. Let \mathcal{P} be the set of points in the plane such that each point in the set is on the boundary of at least two disks. Essentially, this is the set of points at unit distance to centers of at least two disks. We place a vertex at each point in \mathcal{P} and regard the curve between a pair of vertices as an edge, then the embedding of unit disks $\mathcal{B} = \{B_1, \dots, B_n\}$ in the plane gives rise to the drawing P_I of a planar multigraph. Furthermore let D_I be the planar dual of P_I ; it is well known that D_I is also planar.

Next we define a notion of *region* which is essential for the definition of P_G . We call a face \mathcal{R} of the plane graph P_I a *region*, if there exists a nonempty subset $\mathcal{B}' \subseteq \mathcal{B}$ of unit disks such that every point in \mathcal{R} is an interior point of each disk in \mathcal{B}' . Hence with every region \mathcal{R} we can associate a set of unit disks. Since the vertices of G correspond to disks of \mathcal{B} , we can associate a subset of vertices of G , say $\mathcal{V}(\mathcal{R})$, to a region \mathcal{R} . We remark that there could be two regions \mathcal{R}_1 and \mathcal{R}_2 with $\mathcal{V}(\mathcal{R}_1) = \mathcal{V}(\mathcal{R}_2)$. Now we are ready to define the graph P_G . Let $\mathcal{R}_1, \dots, \mathcal{R}_p$ be the regions of P_I . These are faces in P_I and hence in the dual graph D_I we have vertices corresponding to them. That is, in D_I for every region \mathcal{R}_i we have a vertex $v(\mathcal{R}_i)$. We define

$$P_G := D_I[\{v(\mathcal{R}_i) \mid 1 \leq i \leq p\}].$$

Thus, P_G is an induced subgraph of D_I obtained by removing non-region vertices. Figure 1 illustrates the construction of graphs P_I and P_G from unit disks drawing. Next we prove some properties of P_G .

Lemma 1. *Let G be a unit disk graph of maximum degree Δ . Then P_G is a planar graph and every vertex $v \in V(G)$ is a part of at most $3(\Delta^2 + \Delta)$ regions.*

Proof. The graph P_G is a subgraph of D_I , the planar dual of P_I , and hence it is also planar. Let $v \in V(G)$ be a vertex. We consider the embedding (drawing) of unit disks corresponding to the vertices of the closed neighborhood $N_G[v]$ in the plane. Then $|N_G[v]| \leq \Delta + 1$. Let \mathcal{L} be the set of the points in the plane such that each point in the set is on the boundary of at least two disks with distinct center points. This is the induced subgraph of P_I formed by the intersection points of the boundaries of disks from $N_G[v]$. Since every two circles with distinct center points intersect in at most two points, we have that $|\mathcal{L}| \leq 2 \binom{|N_G[v]|}{2} \leq 2 \binom{\Delta+1}{2} = \Delta^2 + \Delta$. Consider the planar graph $P_I[\mathcal{L}]$, which is a subgraph of P_I induced by \mathcal{L} . Observe that v can only be a part of regions defined by faces of $P_I[\mathcal{L}]$. To obtain an upper bound on the number of faces of $P_I[\mathcal{L}]$, we first obtain an upper bound on the number of edges of $P_I[\mathcal{L}]$. First observe that between any pair of vertices in $P_I[\mathcal{L}]$ there can at most be two edges and

there are at most $|\mathcal{L}|$ pairs that have two edges between them. It is well known that a planar graph on n vertices without any parallel edges has at most $3n - 6$ edges. Thus, the number of edges in $P_I[\mathcal{L}]$ is at most $3|\mathcal{L}| - 6 + |\mathcal{L}| = 4|\mathcal{L}| - 6$. Now by Euler's formula, the number of faces in $P_I[\mathcal{L}]$ is at most

$$2 + |E(P_I[\mathcal{L}])| - |\mathcal{L}| \leq 2 + 3|\mathcal{L}| - 6 \leq 2 + 3(\Delta^2 + \Delta) - 6 \leq 3(\Delta^2 + \Delta).$$

Thus v is a part of at most $3(\Delta^2 + \Delta)$ regions. We complete the proof by making a remark that the Euler's formula also holds for graphs with multiple edges. \square

Lemma 2. *Let G be a unit disk graph. Then $\text{tw}(G) \leq (\Delta(G) + 1) \cdot (\text{tw}(P_G) + 1) - 1$.*

Proof. Let $\Delta(G) = \Delta$ and (\mathcal{X}', T) be a tree decomposition of P_G of width $\text{tw}(P_G)$. We build a tree decomposition (\mathcal{X}, T) of G from the tree-decomposition (\mathcal{X}', T) of P_G . Let X'_i be the subset of $V(P_G)$ associated with the node i of T . We define $X_i := \bigcup_{v(\mathcal{R}) \in X'_i} \mathcal{V}(\mathcal{R})$. Recall that $\mathcal{V}(\mathcal{R})$ is a subset of vertices in $V(G)$ characterizing \mathcal{R} . This concludes the description of a decomposition for G . Observe that the set $\mathcal{V}(\mathcal{R})$ is contained in $N_G[w]$ for every $w \in \mathcal{V}(\mathcal{R})$ and hence the size of each of them is bounded above by $\Delta + 1$. Hence the size of each of X_i is at most $(\Delta + 1) \cdot |X'_i|$. This implies that the size of every bag X_i is at most $(\Delta + 1)(\text{tw}(P_G) + 1)$.

Now we show that this is indeed a tree-decomposition for G by proving that it satisfies the three properties of a tree decomposition. By construction, every vertex of $V(G)$ is contained in some X_i . To show that for every edge $uv \in E(G)$ there is a node i such that $u, v \in X_i$, we argue as follows. If there is an edge between u and v in G then unit disks corresponding to these vertices intersect and hence there is a region \mathcal{R} which is completely contained inside this intersection. This implies that $u, v \in \mathcal{V}(\mathcal{R})$. For node i such that $v(\mathcal{R})$ is contained inside X'_i , we have that the corresponding bag X_i contains u and v . To conclude we need to show that for each $v \in V(G)$ the set $Z = \{i \mid v \in X_i\}$ induces a subtree of T . Observe that v appears in all the bags corresponding to node i such that X'_i contains a vertex corresponding to a region which v is a part of. This implies that all these regions are inside the unit disk corresponding to v . Hence the graph induced by vertices corresponding to these regions is connected. Thus the set Z induces a subtree of T . \square

We now show a linear excluded grid theorem for unit disk graphs of bounded degree.

Lemma 3. *Any unit disk graph G with maximum degree Δ contains a $\frac{\text{tw}(G)}{100\Delta^3} \times \frac{\text{tw}(G)}{100\Delta^3}$ grid as a minor.*

Proof. Let G be a unit disk graph of maximum degree Δ , and define P_G as above. Since P_G is planar, by the excluded grid theorem for planar graphs [38], P_G contains a $t \times t$ grid as a minor where $t = \frac{\text{tw}(P_G)}{6}$. By Proposition 1, we know that there is a minor model of this grid, say $\{S[i, j] : 1 \leq i, j \leq t\}$. We know that for every i, j , $P_G[S[i, j]]$ is connected, the sets $S[i, j]$ are pairwise disjoint and finally for every i, j, i', j' such that $|i - i'| + |j - j'| = 1$ there is an edge in P_G with one endpoint in $S[i, j]$ and the other in $S[i', j']$. For every i, j , we build $S'[i, j]$ from $S[i, j]$ by replacing every vertex $v(\mathcal{R}) \in S[i, j]$ by $\mathcal{V}(\mathcal{R})$ and removing duplicates. We set $\Delta' = 3(\Delta^2 + \Delta)$, and observe that for any vertex v in G , Lemma 1 implies that there are at most Δ' sets $S'[i, j]$ that contain v .

We say that an integer pair (i, j) is *internal* if $\Delta' \leq i \leq t - \Delta'$ and $\Delta' \leq j \leq t - \Delta'$. We prove that for any two internal pairs (i, j) and (i', j') such that $|i - i'| + |j - j'| > \Delta'$ the sets $S'[i, j]$ and $S'[i', j']$ are disjoint. To obtain a contradiction assume that both sets contain a vertex v in G . Let X_v be the set of vertices $v(\mathcal{R})$ such that $v \in \mathcal{V}(\mathcal{R})$. We will show that $|X_v| > \Delta'$ which contradicts that v is part of at most Δ' regions. On one hand, $P_G[X_v]$ is connected. On the other hand, both $S'[i, j] \cap X_v$ and $S'[i', j'] \cap X_v$ are non-empty. But any path in P_G between a vertex in $S'[i, j]$ and a vertex in $S'[i', j']$ must pass through at least $\Delta' + 1$ cycles of the grid minor and thus the length of a shortest path between a pair of vertices, $x \in (S'[i, j] \cap X_v)$ and $y \in (S'[i', j'] \cap X_v)$, is at least $\Delta' + 1$. This implies that the length of a shortest path between x and y in $P_G[X_v]$ is at least $\Delta' + 1$ and hence $|X_v| > \Delta'$, yielding the desired contradiction. By an identical argument one can show that, for any two internal pairs (i, j) and (i', j') such that $|i - i'| + |j - j'| > 2\Delta'$ there is no edge with one endpoint in $S'[i, j]$ and the other in $S'[i', j']$.

For every pair a, b of non-negative integers such that $4\Delta'a + 2\Delta' \leq t$ and $4\Delta'b + 2\Delta' \leq t$ we define the sets

- $V[a, b] = \bigcup_{i=0}^{2\Delta-1} \bigcup_{j=0}^{2\Delta-1} S[\Delta + 4\Delta a + i, \Delta + 4\Delta b + j]$.
- $E_h[a, b] = \bigcup_{i=0}^{2\Delta-1} S[3\Delta + 4\Delta a + i, 2\Delta + 4\Delta b]$.
- $E_v[a, b] = \bigcup_{j=0}^{2\Delta-1} S[2\Delta + 4\Delta a, 3\Delta + 4\Delta b + j]$.

One can think of each set $V[a, b]$ as a vertex of a grid, with each set $E_h[a, b]$ being a horizontal edge and each set $E_v[a, b]$ being a vertical edge in this grid. Build $V'[a, b]$ from $V[a, b]$ by replacing every vertex $v(\mathcal{R}) \in V[a, b]$ by $\mathcal{V}(\mathcal{R})$ and removing duplicates. Construct $E'_h[a, b]$ from $E_h[a, b]$ and $E'_v[a, b]$ from $E_v[a, b]$ similarly. We list the properties of the sets $V'[a, b]$, $E'_h[a, b]$ and $E'_v[a, b]$.

1. For every a, b , $G[V'[a, b]]$, $G[E'_h[a, b]]$ and $G[E'_v[a, b]]$ are connected.
2. Distinct sets $V'[a, b]$ and $V'[a', b']$ are pairwise disjoint, and there is no edge with one endpoint in $V'[a, b]$ and the other in $V'[a', b']$.
3. For every a, b the set $E'_h[a, b]$ is disjoint from every set $E'_h[a', b']$, $E'_v[a', b']$ and $V'[a', b']$, except possibly for $V'[a, b]$ and $V'[a + 1, b]$.
4. For every a, b the set $E'_v[a, b]$ is disjoint from every set $E'_h[a', b']$, $E'_v[a', b']$ and $V'[a', b']$, except possibly for $V'[a, b]$ and $V'[a, b + 1]$.
5. For every a, b there is a vertex in $E'_h[a, b]$ which is adjacent to $V'[a, b]$ and a vertex which is adjacent to $V'[a + 1, b]$. Furthermore there is a vertex in $E'_v[a, b]$ which is adjacent to $V'[a, b]$ and a vertex which is adjacent to $V'[a, b + 1]$.

Property 1 follows directly from the fact that $P_G[V[a, b]]$, $P_G[E_h[a, b]]$ and $P_G[E_v[a, b]]$ are connected. Properties 2, 3 and 4 follow from the fact that for any two internal pairs (i, j) and (i', j') such that $|i - i'| + |j - j'| > 2\Delta'$ the sets $S[i, j]$ and $S[i', j']$ are disjoint and have no edges between each other. Finally, Property 5 follows from the fact that for every a, b there is a vertex in $E_h[a, b]$ which is adjacent to $V[a, b]$ and a vertex which is adjacent to $V[a + 1, b]$, and that there is a vertex in $E_v[a, b]$ which is adjacent to $V[a, b]$ and a vertex which is adjacent to $V[a, b + 1]$.

For a pair a, b of integers consider the set $E'_h[a, b]$. The properties 1, 2 and 5 ensure that some connected component $E_h^*[a, b]$ of $G[E'_h[a, b] \setminus (V'[a, b] \cup V'[a + 1, b])]$ contains at least one neighbour of $V'[a, b]$ and one neighbour of $V'[a + 1, b]$. Similarly at least one connected component $E_v^*[a, b]$ of $G[E'_v[a, b] \setminus (V'[a, b] \cup V'[a, b + 1])]$ contains at least one neighbour of $V'[a, b]$ and one neighbour of $V'[a, b + 1]$. Then the family

$$\{V'[a, b], E_h^*[a, b], E_v^*[a, b] : 4\Delta'a + 2\Delta' \leq t \text{ and } 4\Delta'b + 2\Delta' \leq t\}$$

of vertex sets in G forms a model of a $\lfloor \frac{t-2\Delta'}{4\Delta'} \rfloor \times \lfloor \frac{t-2\Delta'}{4\Delta'} \rfloor$ grid minor in G with every edge subdivided once. The sets $V'[a, b]$ are models of the vertices of the grid, the sets $E_h^*[a, b]$ are models of the subdivision vertices on the horizontal edges, while $E_v^*[a, b]$ are models of the subdivision vertices on the vertical edges. Now by Lemma 2, we know that $\text{tw}(P_G) \geq \frac{\text{tw}(G)+1}{(\Delta+1)} - 1$. Combining this with the fact that $t = \frac{\text{tw}(P_G)}{6}$ we can show that G has a grid of size $\frac{\text{tw}(G)}{100\Delta^3} \times \frac{\text{tw}(G)}{100\Delta^3}$ as a minor. This concludes the proof. \square

Using Lemma 3 we show the following theorem and then combining this theorem with Observation 1 we get an analogous result for K_t -free unit disk graphs.

Theorem 1. *Let \mathcal{G}_U^Δ be the class of unit disk graphs such that the maximum degree of every graph G in \mathcal{G}_U^Δ is at most Δ . Then \mathcal{G}_U^Δ has truly sublinear treewidth with $\lambda = \frac{1}{2}$.*

Proof. Let $G \in \mathcal{G}_U^\Delta$ have a vertex set X such that $\text{tw}(G \setminus X) \leq \eta$. Suppose for contradiction that $\text{tw}(G) \geq (\eta + 1)200\Delta^3\sqrt{|X|}$. Then, by Lemma 3, G contains a $(2(\eta + 1)\sqrt{|X|}) \times (2(\eta + 1)\sqrt{|X|})$ grid as a minor. This grid contains $4|X|$ vertex disjoint $(\eta + 1) \times (\eta + 1)$ grids. However, the set X must intersect each of these grids, as $\text{tw}(G \setminus X) \leq \eta$ and the treewidth of each grid is $\eta + 1$. But then $|X| \geq 4|X|$, a contradiction. \square

Corollary 1. *Let \mathcal{G}_U^t be the class of unit disk graphs such that every graph G in \mathcal{G}_U^t is K_t -free. Then \mathcal{G}_U^t has truly sublinear treewidth with $\lambda = \frac{1}{2}$.*

3.2 Structure of K_t -free Map Graphs.

In this section we show that map graphs with bounded clique size have truly sublinear treewidth. It is known that for every map graph G one can associate a map $\mathcal{M} = (\mathcal{E}, \omega)$ such that (i) no vertex in \mathcal{E} is incident only to lakes; (ii) there are no edges in \mathcal{E} whose two incident faces are both lakes (possibly the same lake); (iii) every vertex in \mathcal{E} is incident to at most one lake, and incident to such a lake at most once [20, p. 149]. From now onwards we will assume that we are given map satisfying the above properties. For our proof we also need the following combinatorial lemma.

Lemma 4. *Let G be a map graph associated with \mathcal{M} such that the maximum clique size in G is at most t . Then the maximum vertex degree of \mathcal{E} is at most $t + 2$.*

Proof. Targeting towards a contradiction, let us assume that there is a vertex $v \in V(\mathcal{E})$ of degree at least $t + 3$. By definition, each connected component of \mathcal{E} is biconnected and hence there are at least $t + 2$ cyclic faces adjacent to v . By the properties of \mathcal{M} , we have that all, except maybe one, adjacent faces are not lakes. However the vertices corresponding to nation faces form a clique of size $t + 1$ in G , a contradiction. \square

For our proof we also need the notions of radial and dual of map graphs. The *radial graph* $R = R(\mathcal{M})$ has a vertex for every vertex of \mathcal{E} and for every nation of \mathcal{E} , and R is a bipartite graph with bipartition $V(\mathcal{E})$ and $N(\mathcal{E})$. Two vertices $v \in V(\mathcal{E})$ and $f \in N(\mathcal{E})$ are adjacent in R if v is incident to nation f . The *dual* $D = D(\mathcal{M})$ of \mathcal{M} has vertices corresponding only to the nations of \mathcal{E} . The graph D has a vertex for every nation of \mathcal{E} , and two vertices are adjacent in D if the corresponding nations of G share an edge. We now show a linear excluded grid theorem for map graphs with bounded maximum clique.

Lemma 5. *There exists a constant ρ such that any map graph G with maximum clique size t contains a $\frac{\rho \cdot \text{tw}(G)}{t} \times \frac{\rho \cdot \text{tw}(G)}{t}$ grid minor.*

Proof. Let \mathcal{M} be the map such that the graph associated with it is G . We now apply the result from [20, Lemma 4] that states that the treewidth of the map graph G is at most the product of the maximum vertex degree in \mathcal{E} and $\text{tw}(R) + 1$. By Lemma 4, we know that the maximum vertex degree of \mathcal{E} is at most $t + 2$, and hence $\text{tw}(G) \leq (t + 2) \cdot (\text{tw}(R) + 1)$. We now apply [20, Lemma 3] which bounds the treewidth of a radial graph of a map. In particular, by [20, Lemma 3] we have that for R , the radial graph of \mathcal{M} , $\text{tw}(R) = O(\text{tw}(D))$. This implies that $\text{tw}(G) = O(t \cdot \text{tw}(D))$. Observe that the graph D , the dual of \mathcal{M} , is a planar subgraph of G . By a result of Robertson et al. [38], we have that for every planar graph H there exists a constant d such that it has $d \cdot \text{tw}(H) \times d \cdot \text{tw}(H)$ grid graph as a minor. This implies that there exists a constant d such that D has $d \cdot \text{tw}(D) \times d \cdot \text{tw}(D)$ grid graph as a minor. This combined with facts that $\text{tw}(G) = O(t \cdot \text{tw}(D))$ and D is a subgraph of G implies that there exists a constant ρ such that G has $\frac{\rho \cdot \text{tw}(D)}{t} \times \frac{\rho \cdot \text{tw}(D)}{t}$ grid graph as a minor. \square

Theorem 2. *Let \mathcal{G}_M^t be the class of map graphs such that the every graph $G \in \mathcal{G}_M^t$ is K_t -free. Then \mathcal{G}_M^t has truly sublinear treewidth with $\lambda = \frac{1}{2}$.*




Figure 2: Family of disks used to construct graph G_t for $t = 3$

Proof. Let $G \in \mathcal{G}_M^t$ have a vertex set X such that $\text{tw}(G \setminus X) \leq \eta$. Suppose for contradiction that $\text{tw}(G) \geq (\eta + 1)2\rho t\sqrt{|X|}$. Then, by Lemma 5, G contains a $(2(\eta + 1)\sqrt{|X|}) \times (2(\eta + 1)\sqrt{|X|})$ grid as a minor. This grid contains $4|X|$ vertex disjoint $(\eta + 1) \times (\eta + 1)$ grids. However, the set X must intersect each of these grids, as $\text{tw}(G \setminus X) \leq \eta$ and the treewidth of each grid is $\eta + 1$. But then $|X| \geq 4|X|$, a contradiction. \square

3.3 K_4 -Free Disc Graphs.

Our result in Section 3.1 can easily be generalized to disk graphs of bounded degree. There is another widely used concept of *ply* related to geometric graphs which has turned out to be very useful algorithmically [41]. An intersection graph G generated by set of disks $\mathcal{B} = \{B_1, \dots, B_n\}$ (not necessarily unit disks) is said to have ply ℓ if every point in the plane is contained inside at most ℓ disks in \mathcal{B} . Observe that if a unit disk graph has bounded ply then it also has bounded vertex degree but this is not true for disk graphs. Here, we show that already disk graphs with ply 3 and K_4 -free disk graphs do not have truly sublinear treewidth.

Theorem 3. *K_4 -free disk graphs and disk graphs with ply 3 do not have truly sublinear treewidth.*

For the proof of Theorem 3, we need the concept of a *bramble*. A *bramble* in a graph G is a family of connected subgraphs of G such that any two of these subgraphs have a nonempty intersection or are joined by an edge. The *order* of a bramble is the minimum number of vertices required to hit all subgraphs in the bramble. Seymour and Thomas [39] proved that a graph has treewidth k if and only if the maximum order of a bramble of G is $k + 1$. Thus a bramble of order $k + 1$ is a witness that the graph has treewidth at least k . We will use this characterization to get a lower bound on the treewidth of the graph we construct.

Proof. We define a family \mathcal{F} of disk graphs of ply 3 such that for every $G \in \mathcal{F}$ we can find a set $X \subseteq V(G)$ such that $\text{tw}(G \setminus X) \leq 1$ while $\text{tw}(G) \geq |X| - 1$. Given a natural number $t \geq 2$, our graph G_t is defined as follows. We give the coordinates for centers of these disks.

- We have “small” disks of radius 0.99 centered at $(1.25p, 2q)$ for $0 \leq p \leq 3t^2$ and $0 \leq q \leq t - 1$.
- We have “large” disks with radius $t - 0.01$ centered at $((2p + 1)t, t)$, $0 \leq p \leq t - 1$.

Intuitively, we have small disks stacked in t rows, where in each row two consecutive disks intersect. Large disks intersect some unit disks in each row and they are pairwise disjoint among themselves. See Figure 2, for an example of our construction. Let G_t be the disk graph obtained from the intersection of disks placed as above. Observe that every point in the plane only occurs in at most 3 disks and hence the ply of the graph is 3. Furthermore, since at most 3 disks mutually intersect we have that G is also K_4 -free. Let A be the set of vertices corresponding to small disks in rows and X be the set of remaining vertices. Observe that the graph induced by A is a set of vertex disjoint paths and hence $\text{tw}(G_t[A]) = \text{tw}(G_t \setminus X) = 1$. We show that the treewidth of G_t is at least $t - 1$ by exhibiting a bramble of order t . Let us take the following set S_i , $0 \leq i \leq t - 1$. The set S_i consists of vertices

corresponding to small disks centered at $(1.25p, 2i)$, where $0 \leq p \leq 3t^2$ and a vertex corresponding to large disk with radius $t - 0.01$ centered at $((2i + 1)t, t)$. Since the disk with radius $t - 0.01$ intersects at least one small disk in each row we have that the sets S_i mutually intersect. Furthermore $S_i \cap S_j = \emptyset$ for all $i \neq j$. This implies that the smallest number of vertices required to cover all S_i is at least t . This implies that $\text{tw}(G) \geq t - 1 = |X| - 1$. \square

4 Applications

In this section we show that every reducible minor-bidimensional problem with the separation property has EPTAS on unit disk graphs and map graphs with bounded maximum clique. Finally, we will show how we can extend these results to unit disk graphs and map graphs for variety of problems. We recall that a PTAS is an algorithm which takes an instance I of an optimization problem and a parameter $\epsilon > 0$ and, runs in time $n^{\mathcal{O}(f(1/\epsilon))}$, produces a solution that is within a factor ϵ of being optimal. A PTAS with running time $f(1/\epsilon) \cdot n^{\mathcal{O}(1)}$, is called an efficient PTAS (EPTAS).

EPTAS on \mathcal{G}_U^t and \mathcal{G}_M^t . Towards our goal we need the following result from [26].

Proposition 2 ([26]). *Let Π be an η -transversable, reducible graph optimization problem. Then Π has an EPTAS on every graph class \mathcal{G} with truly sublinear treewidth.*

To use this result we only need to show that every reducible minor-bidimensional problem with the separation property is η -transversable for some η on \mathcal{G}_U^t and \mathcal{G}_M^t . For every fixed integer η we define the η -TRANSVERSAL problem as follows. Input is a graph G , and the objective is to find a minimum cardinality vertex set $S \subseteq V(G)$ such that $\text{tw}(G \setminus S) \leq \eta$. We now give a polynomial time constant factor approximation for the η -TRANSVERSAL problem on \mathcal{G}_U^t and \mathcal{G}_M^t . The proof of the following result is similar to the one in [26, Lemma 4.1]. We give the proof here for completeness.

Lemma 6. *For every integer η there is a constant c and a polynomial time c -approximation algorithm for the η -TRANSVERSAL problem on \mathcal{G}_U^t and \mathcal{G}_M^t .*

For the proof of Lemma 6 we also need the following result from [26, Lemma 3.2].

Lemma 7. *Let \mathcal{G} be a hereditary graph class of truly sublinear treewidth with parameter λ . For any $\epsilon < 1$ there is a γ such that for any $G \in \mathcal{G}$ and $X \subseteq V(G)$ with $\text{tw}(G \setminus X) \leq \eta$ there is a $X' \subseteq V(G)$ satisfying $|X'| \leq \epsilon|X|$ and for every connected component C of $G \setminus X'$ we have $|C \cap X| \leq \gamma$ and $|N(C)| \leq \gamma$. Moreover X' can be computed from G and X in polynomial time, where the polynomial is independent of ϵ , λ and η .*

Proof of Lemma 6. Let X be a smallest vertex set in G such that $\text{tw}(G \setminus X) \leq \eta$. By Lemma 7 with $\epsilon = 1/2$ there exists a γ depending only on t and η and a set X' with $|X'| \leq |X|/2$ such that for any component C of $G \setminus X'$, $|C \cap X| \leq \gamma$ and $|N(C)| \leq \gamma$. Since X is the *smallest* set such that $\text{tw}(G \setminus X) \leq \eta$, there is a component C of $G \setminus X'$ with treewidth strictly more than η . Let $Z = N(C)$ and observe that $Z \subseteq X'$ is a set of size at most γ such that C is a connected component of $G \setminus Z$.

The algorithm proceeds as follows. It tries all possibilities for Z and looks for a connected component C of $G \setminus Z$ such that $\eta < \text{tw}(G[C]) = O(\sqrt{\gamma})$. It solves the η -TRANSVERSAL problem optimally on $G[C]$ by noting that η -TRANSVERSAL can be formulated as a MIN-CMSO problem and applying the algorithm by Borie et al. [8]. Let X_C be the solution obtained for $G[C]$. The algorithm adds X_C and $N(C)$ to the solution and repeats this step on $G \setminus (C \cup N(C))$ as long as $\text{tw}(G) \geq \eta$.

Clearly, the set returned by the algorithm is a feasible solution. We now argue that the algorithm is a $(\gamma + 1)$ -approximation algorithm. Let C_1, C_2, \dots, C_t be the components found by the algorithm in this manner. Since X must contain at least one vertex in each C_i it follows that $t \leq |X|$. Now, for each i , $N(C_i)$ contains at most γ vertices outside of $\bigcup_{j < i} N(C_j)$. Thus $\bigcup_{i \leq t} N(C_i) \leq \gamma|X|$. Furthermore for

each C , $|X_C| \leq |X \cap C|$ and hence the size of the returned solution is at most $(\gamma + 1)|X|$, which proves the lemma. \square

By exchanging the Excluded grid theorem for H -minor free graphs by Linear Grid lemmas (Lemmas 3 and 5), we can adapt the proof of [27, Lemma 3.2] to show the following.

Lemma 8. *Let Π be a minor-bidimensional problem with the separation property. There exists a constant η such that for every G in \mathcal{G}_U^t or \mathcal{G}_M^t , there is a subset $S \subseteq V(G)$ such that $|X| = O(t^{O(1)}\pi(G))$, and $\text{tw}(G \setminus S) \leq \eta$.*

For the proof of the Lemma 8, we also need the following well known lemma, see e.g. [7], on separators in graphs of bounded treewidth.

Lemma 9. *Let G be a graph of treewidth at most t and $w : V(G) \rightarrow \{0, 1\}$ be a weight function. Then there is a partition of $V(G)$ into $L \uplus S \uplus R$ such that*

- $|S| \leq t + 1$, $N(L) \subseteq S$ and $N(R) \subseteq S$,
- every connected component $G[C]$ of $G \setminus S$ has $w(C) \leq w(V)/2$,
- $\frac{w(V(G)) - w(S)}{3} \leq w(L) \leq \frac{2(w(V(G)) - w(S))}{3}$ and $\frac{w(V(G)) - w(S)}{3} \leq w(R) \leq \frac{2(w(V(G)) - w(S))}{3}$.

Proof of Lemma 8. We prove our result for graphs in \mathcal{G}_U^t . The result for graphs in \mathcal{G}_M^t is analogous. Since Π is a minor-bidimensional problem, there exists a constant $\delta > 0$ such that $\pi(R) \geq \delta r^2$ on $r \times r$ grid minor. The definition of minor-bidimensionality together with Lemma 3 imply that for every $G \in \mathcal{G}_U^t$ we have that $\text{tw}(G) \leq 10t^3 \sqrt{\frac{\pi(G)}{\delta}}$. Thus, there is a constant d' depending on t and δ such that $\text{tw}(G) \leq d' \sqrt{\pi(G)}$.

We first make the following observation. For any numbers $a > 0$, $b > 0$, since $\lambda < 1$, we have that $a^\lambda + b^\lambda > (a + b)^\lambda$. Thus we have $\rho = \min_{1/3 \leq \alpha \leq 2/3} \alpha^\lambda + (1 - \alpha)^\lambda > 1$. Now we proceed to our proof.

We want to construct a set X such that the treewidth of $G[V(G) \setminus X]$ is at most η (to be fixed later). Let us fix a solution Z of size $\pi(G)$ for G and a weight function $w : V(G) \rightarrow \{0, 1\}$ that assigns 1 to every vertex in Z and 0 otherwise. By Lemma 9, there is a partition of $V(G)$ into L , S and R such that $|S| \leq d' \sqrt{\pi(G)} + 1$, $N(L) \subseteq S$, $N(R) \subseteq S$, $|L \cap Z| \leq 2\pi(G)/3$ and $|R \cap Z| \leq 2\pi(G)/3$. By deleting S from the graph G , we obtain two graphs $G[L]$ and $G[R]$ with no edges between them. Since Π is separable, there exists a constant β such that $\pi(G[L]) \leq |Z \cap L| + \beta|S|$ and $\pi(G[R]) \leq |Z \cap R| + \beta|S|$. Thus we put S into X and then proceed recursively in $G[L]$ and $G[R]$. Since Π is minor-bidimensional problem with the separation property, we have that in recursive step for a graph G' with solution of size ℓ we find a separator of size $\mathcal{O}(\sqrt{\ell})$. We set $k = \pi(G)$. Then the size of the set X we are looking for is governed by the following recurrence.

$$T(k) \leq \max_{\frac{1}{3} \leq \alpha \leq \frac{2}{3}} \left\{ T(\alpha k + \beta\sqrt{k}) + T((1 - \alpha)k + \beta\sqrt{k}) + d'\sqrt{k} + 1 \right\}. \quad (1)$$

We first set two constants q and γ which will be used in the proof of above recurrence. Set $q = \frac{6\beta + 3d'}{\rho - 1}$ and $\gamma = \lceil 4\delta^{-1} \rceil + (3\beta)^2 + (3q)^2 + 1$. The base case of the recursion is when $k \leq \gamma$ and once we reach this case, we do not decompose the graph any further. Thus for the base case we set $T(k) = 0$. Now using induction we can show that the size of $|X| \leq k - q\sqrt{k}$ and for every component C in $G[V(G) \setminus X]$ we have that $\pi(C) = O(\gamma)$. We first show that if $k \geq \gamma/3$ then $T(k) \leq k - q\sqrt{k}$ by induction on k . For the base case if $\gamma/3 \leq k \leq \gamma$ then the choice of γ implies that $k - q\sqrt{k} \geq \frac{\gamma}{3} - q\sqrt{\gamma} \geq 0 = T(k)$.

We now consider $T(k)$ for $k > \gamma$. By our choice of q we have that for all $k > \gamma$, $\alpha k + \beta\sqrt{k} < k$ and $(1 - \alpha)k + \beta\sqrt{k} < k$. The induction hypothesis then yields the following inequality.

$$\begin{aligned}
T(k) &\leq \max_{1/3 \leq \alpha \leq 2/3} T(\alpha k + \beta\sqrt{k} + 1) + T((1 - \alpha)k + \beta\sqrt{k} + 1) + d'\sqrt{k} + 1 \\
&\leq \max_{1/3 \leq \alpha \leq 2/3} k - q\sqrt{\alpha k} - q(\sqrt{(1 - \alpha)k}) + 2(\beta\sqrt{k} + 1) + d'\sqrt{k} + 1 \\
&\leq \max_{1/3 \leq \alpha \leq 2/3} k - q\sqrt{k}(\alpha^{1/2} + (1 - \alpha)^{1/2}) + 2(\beta\sqrt{k} + 1) + d'\sqrt{k} + 1 \\
&\leq k - q\sqrt{k} - q(\rho - 1)\sqrt{k} + 2(\beta\sqrt{k} + 1) + d'\sqrt{k} + 1 \\
&\leq k - q\sqrt{k}.
\end{aligned}$$

The last inequality holds whenever $q(\rho - 1)\sqrt{k} \geq 2(\beta\sqrt{k} + 1) + d'\sqrt{k} + 1$, which is ensured by the choice of q and the fact that $\sqrt{k} \geq 1$. Thus $T(k) \leq k$ for all k . This proves the bound on the size of X . Notice that every component C in $G[V(G) \setminus X]$ has at most γ vertices and thus by bidimensionality we have that $\text{tw}(C) = O(\sqrt{\gamma})$. We set η as the $\mu\sqrt{\gamma}$, where μ is the constant appearing in the term $O(\sqrt{\gamma})$. This proves the theorem. \square

Theorem 4. *Let Π be a reducible minor-bidimensional problem with the separation property. There is an EPTAS for Π on \mathcal{G}_U^t and \mathcal{G}_M^t .*

Proof. Combining Lemmas 6 and 8 we get that every reducible minor-bidimensional problem Π with the separation property is η -transversable for some η on \mathcal{G}_U^t and \mathcal{G}_M^t . Thus the theorem follows by applying Proposition 2 in combination with Theorem 2 and Corollary 1. \square

EPTAS on Unit Disc Graphs and Map Graphs. In this section we give EPTAS for several problems on unit disk graphs and map graphs. Our first problem is the following generic problem. Let \mathcal{F} be a finite set of graphs. In the \mathcal{F} -DELETION problem, we are given an n -vertex graph G as an input, and asked to find a minimum sized subset $S \subseteq V(G)$ such that $G \setminus S$ does not contain a graph from \mathcal{F} as a minor. We refer to such a subset S as a \mathcal{F} -hitting set. The \mathcal{F} -DELETION problem is a generalization of several fundamental problems. For example, when $\mathcal{F} = \{K_2\}$, a complete graph on two vertices, this is the VERTEX COVER problem. When $\mathcal{F} = \{C_3\}$, a cycle on three vertices, this is the FEEDBACK VERTEX SET problem. Other famous cases are $\mathcal{F} = \{K_{2,3}, K_4\}$, $\mathcal{F} = \{K_{3,3}, K_5\}$ and $\mathcal{F} = \{K_3, T_2\}$, which correspond to removing vertices to obtain outerplanar graphs, planar graphs, and graphs of pathwidth one respectively. Here, $K_{i,j}$ denotes the complete bipartite graph with bipartitions of sizes i and j , and K_i denotes the complete graph on i vertices. Further, a T_2 is a star on three leaves, each of whose edges has been subdivided exactly once. In literature, these problems are known as OUTERPLANAR DELETION SET, PLANAR DELETION SET and PATHWIDTH ONE DELETION SET respectively. Now we show that if \mathcal{F} contains a planar graph then \mathcal{F} -DELETION has EPTAS on unit disk graphs and map graphs. It is known from [26] that \mathcal{F} -DELETION problem is reducible minor-bidimensional problem with the separation property whenever \mathcal{F} contains a planar graph.

Theorem 5. *Let \mathcal{F} be a finite set of graphs containing a planar graph. Then \mathcal{F} -DELETION admits an EPTAS on unit disk graphs and map graphs.*

Proof. Let G be the input graph, ϵ be a fixed constant and \mathcal{F} be an obstruction set containing a planar graph of size h . This implies that any optimal \mathcal{F} -hitting set in G must contain all but $h - 1$ vertices from any clique in G . We outline a proof below only for unit disk graphs, the proof for map graphs is similar.

The algorithm proceeds as follows. It finds a maximum clique C of G . One can find a maximum sized clique in unit disk graphs and map graphs in polynomial time [11, 12, 13, 37]. The algorithm adds C to the solution and repeats this step on $G \setminus C$ as long as there is a clique of size $\frac{(1+\epsilon)h}{\epsilon}$. Once we have that the maximum size of a clique is bounded by $\frac{(1+\epsilon)h}{\epsilon}$, we can use the EPTAS obtained in Theorem 4 to get a \mathcal{F} -hitting set of G of size $(1 + \epsilon)OPT$, where OPT is the size of a minimum \mathcal{F} -hitting set.

Clearly, the set returned by the algorithm is a feasible solution. We now argue that the algorithm is an EPTAS. Clearly the running time of our algorithm is of desired form. As the step where we find a clique and add all its vertices in our solution can be done in polynomial time and finally we run an EPTAS on a graph where the maximum degree is bounded by a function of ϵ . Let X be an optimal \mathcal{F} -hitting set of G . Let C_1, C_2, \dots, C_t be the cliques found by the algorithm and G_q be the graph where we apply Theorem 4. Since X must contain at least $|C_i| - h$ vertices and $|C_i| \geq \frac{(1+\epsilon)h}{\epsilon}$, we have that $|C_i| \leq (1+\epsilon)(|C_i| - h) \leq (1+\epsilon)(|X \cap C_i|)$. Thus the size of the solution returned by the algorithm satisfies the following inequality $\sum_{i=1}^t |C_i| + (1+\epsilon)|X \cap V(G_q)| \leq (1+\epsilon)(\sum_{i=1}^t |X \cap C_i| + |X \cap V(G_q)|) \leq (1+\epsilon)|X| = (1+\epsilon)OPT$. This completes the proof. \square

Next we show how we can obtain an EPTAS for CONNECTED VERTEX COVER on unit disk graphs and map graphs. In CONNECTED VERTEX COVER we are given a graph G and the objective is to find a minimum size subset $S \subseteq V(G)$ such that $G[S]$ is connected and every edge in $E(G)$ has at least one endpoint in S .

Theorem 6. CONNECTED VERTEX COVER admits an EPTAS on unit disk graphs and map graphs.

Proof. Observe that CONNECTED VERTEX COVER is 0-transversable. Given a graph G we find a maximal matching in linear time and output the endpoints of the matching as X . Any vertex cover must contain at least one endpoint from each edge in the matching, and thus $|X| \leq 2\pi(G)$. Also, $\text{tw}(G \setminus X) = 0$. We recall the proof from [26] that CONNECTED VERTEX COVER is reducible, as we will use this to get an EPTAS here. Given a graph G and set X , let $G' = G \setminus X$ and let $R = N(X)$. The annotated problem Π' is to find a minimum sized set $S' \subseteq V(G')$ such that every edge in G' has an end point in S' and every connected component of $G'[S']$ contains a vertex in R . Notice that for any connected vertex set S of G , $S \setminus X$ is a feasible solution to Π' on G' . Conversely, for any feasible solution S' of Π' on G' , we have that $S = S' \cup X$ is a vertex cover of G and has at most $|X|$ connected components. Since S is a vertex cover it is sufficient to add $(|X| - 1)$ vertices to S in order to make it a connected vertex cover of G . Hence, CONNECTED VERTEX COVER is reducible. One can similarly show that the annotated problem Π' is 0-transversable and reducible. This implies that Π' has EPTAS on \mathcal{G}_U^t and \mathcal{G}_M^t .

To get our EPTAS for CONNECTED VERTEX COVER we do similar to what we did for \mathcal{F} -DELETION problem in Theorem 5. The only change is that we keep finding clique and including it in our solution until there is no clique of size $\frac{(2+\epsilon)}{\epsilon}$. Let C_1, C_2, \dots, C_q be the cliques found by the algorithm and G_q be the graph on which we apply Proposition 2. Let Z be the union of cliques, that is, $Z = \cup_{i \leq q} C_i$. Now we define the annotated problem Π' with respect to set Z and using Proposition 2 obtain a set W of size $(1+\epsilon)OPT'$, where OPT' is the size of a minimum cardinality set in G_q such that every edge in G_q has an end point in W and every connected component of $G_q[W]$ contains a vertex in $R = N(Z) \cap V(G_q)$. Now consider the set $W \cup Z$. This is a vertex cover of G such that it has q components and hence we can make it connected by adding at most $q - 1$ vertices. Let the final solution returned by our algorithm be S . Let X be an optimal connected vertex cover of G . Since X must contain at least $|C_i| - 1$ vertices and the size of $|C_i| \geq \frac{(2+\epsilon)}{\epsilon}$, we have that $|C_i| + 1 \leq (1+\epsilon)(|C_i| - 1) \leq (1+\epsilon)(|X \cap C_i|)$. Thus the size of the solution returned by the algorithm satisfies the following inequality $\sum_{i=1}^t (|C_i| + 1) + (1+\epsilon)|X \cap V(G_q)| \leq (1+\epsilon)(\sum_{i=1}^t (|X \cap C_i|) + |X \cap V(G_q)|) \leq (1+\epsilon)|X| = (1+\epsilon)OPT$. This completes the proof. \square

EPTAS for CYCLE PACKING on Unit Disk Graphs. A cycle packing in a graph G is a collection C_1, C_2, \dots, C_p of pairwise disjoint vertex sets such that for every i , $G[C_i]$ induces a cycle. The integer p is the size of the cycle packing and in the CYCLE PACKING problem the objective is to find a cycle packing of maximum size in the input graph. CYCLE PACKING is known to be minor-bidimensional, separable and reducible [26]. Thus by Theorem 4, the problem admits an EPTAS on \mathcal{G}_U^t . Hence, in order to give an EPTAS for CYCLE PACKING on unit disk graphs it is sufficient to prove the following lemma.

In particular, the following lemma implies that if we find a sufficiently large clique X , partition X into triangles and add this partition to our packing, this will give a good approximation of how the optimum solution intersects with X . Here a *triangle* is a cycle on three vertices.

Lemma 10. *Let G be a unit disk graph and X be a clique in G . There is a maximum size cycle packing C_1, C_2, \dots, C_p in G such that at most 405 cycles C_i in the packing satisfy $C_i \cap X \neq \emptyset$ and $C_i \setminus X \neq \emptyset$.*

Proof. Let X be a clique in G . The centers of all disks corresponding to vertices of X must be inside a 2×2 square. Thus the centers of all disks corresponding to vertices in $N(X)$ must be in a 6×6 square. By [22, Lemma 2], the vertices in $N(X)$ can be partitioned into 27 cliques S_1, S_2, \dots, S_{27} . Note that in the definition of unit disk graphs used in [22, Lemma 2] two vertices are adjacent if the centers of the corresponding disks is at distance at most 1 from each other, while in this paper two vertices are adjacent if the centers of their disks are at distance at most 2. This difference is taken into account when applying [22, Lemma 2]. We say that a cycle C *crosses* X if $C \cap X \neq \emptyset$ and $C \setminus X \neq \emptyset$. Let C_1, C_2, \dots, C_p be a maximum cycle packing in G that has the fewest cycles crossing X . Observe that any cycle C that crosses X intersects with X in at most two vertices — since otherwise $G[C \cap X]$ induces a triangle, say T and then we can replace C by T in the cycle packing and obtain a maximum size cycle packing with fewer cycles that cross X . This contradicts the choice of the packing C_1, C_2, \dots, C_p .

We prove that there can be at most 54 cycles in the packing that intersect X in exactly 2 vertices. Suppose for contradiction that there are at least 55 such cycles. Each such cycle contains at least one vertex in $N(X)$. Since each vertex in $N(X)$ is in one of the 27 cliques S_1, \dots, S_{27} the pigeon hole principle implies that there are three cycles C_a, C_b and C_c in the packing which all intersect X in exactly two vertices and a clique S_i such that $C_a \cap S_i \neq \emptyset, C_b \cap S_i \neq \emptyset$ and $C_c \cap S_i \neq \emptyset$. Since all cycles in the packing are vertex disjoint, this means that $S_i \cap (C_a \cup C_b \cup C_c)$ contains a triangle T_1 . On the other hand, $X \cap (C_a \cup C_b \cup C_c)$ is a clique on 6 vertices, and can be partitioned into two triangles T_2 and T_3 . Now we can remove C_a, C_b and C_c from the proposed packing and replace them by T_1, T_2 and T_3 . The resulting packing has the same size, but fewer cycles that cross X . This contradicts the choice of the packing C_1, C_2, \dots, C_p .

Now we show that there can be at most $2(27 \times 27) = 1458$ cycles in the packing that intersect with X in exactly 1 vertex. Every such cycle contains at least two vertices in $N(X)$. For a pair (i, j) of integers $1 \leq i \leq j \leq 27$ we say that a cycle C_a is an (i, j) cycle if C_a contains two distinct vertices u and v such that $u \in S_i$ and $v \in S_j$. If there are more than 1458 cycles in the packing that intersect with X in exactly 1 vertex then there are i and j such that there are three (i, j) -cycles C_a, C_b and C_c in the packing that intersect X in one vertex. Let u_a, u_b and u_c be three vertices in $C_a \cap S_i, C_b \cap S_i$ and $C_c \cap S_i$ respectively. Similarly, let v_a, v_b and v_c be the three vertices in $C_a \cap S_j, C_b \cap S_j$ and $C_c \cap S_j$ respectively. Now $T_1 = \{u_a, u_b, u_c\}, T_2 = \{v_a, v_b, v_c\}$ and $T_3 = X \cap (C_a \cup C_b \cap C_c)$ are vertex disjoint triangles. We can replace $C_a, C_b,$ and C_c by T_a, T_b and T_c in the cycle packing and obtain a maximum size cycle packing with fewer cycles that cross X , contradicting the choice of C_1, \dots, C_p . Hence there are at most $27 + 1458 = 1485$ cycles in the packing that cross X . \square

Theorem 7. CYCLE PACKING *admits an EPTAS on unit disk graphs.*

Proof. Given a unit disk graph G and ϵ , we choose t to be $\frac{(1485 \times 3) = 4455}{\epsilon}$. If G does not contain a clique of size t then we apply the EPTAS for CYCLE PACKING on \mathcal{G}_U^ϵ guaranteed by Theorem 4 to give a $(1 - \epsilon)$ -approximation for CYCLE PACKING. If G contains a clique X of size t , the algorithm partitions X into $\frac{\lfloor X \rfloor}{3}$ triangles T_1, \dots, T_x , recursively finds a $(1 - \epsilon)$ -approximate cycle packing C_1, \dots, C_p in $G \setminus X$ and returns $T_1, \dots, T_x, C_1, \dots, C_p$ as an approximate solution. Clearly, the algorithm terminates in $f(\epsilon) \cdot n^{O(1)}$ time, so it remains to argue that the returned solution is indeed a $(1 - \epsilon)$ -approximate cycle packing of G . We prove this by induction on the number n of vertices in G . Let OPT be the size of the largest cycle packing in G .

If there is no clique of size t and we apply the EPTAS for CYCLE PACKING on \mathcal{G}_U^ϵ , then clearly the returned solution is a $(1 - \epsilon)$ -approximation. If the algorithm finds such a clique X , Lemma 10 ensures

that there is a cycle packing of size OPT such that at most 1485 cycles in the packing cross X . All cycles in the packing that intersect with X but do not cross X are triangles in X . Hence $G \setminus X$ contains a cycle packing of size at least $OPT - \frac{|X|}{3} - 1485$. By the inductive hypothesis the algorithm returns a cycle packing in $G \setminus X$ of size at least $(OPT - \frac{|X|}{3} - 1485)(1 - \epsilon)$. Now, X contains $\frac{|X|}{3}$ triangles T_1, \dots, T_x . Hence, the total size of the packing returned by the algorithm is at least

$$\left(OPT - \frac{|X|}{3} - 1485\right)(1 - \epsilon) + \frac{|X|}{3} = OPT(1 - \epsilon) - \left(\frac{|X|}{3} + 1485\right)(1 - \epsilon) + \frac{|X|}{3} \geq OPT(1 - \epsilon)$$

since $|X| \geq t$. This concludes the proof. \square

EPTAS for (CONNECTED) VERTEX COVER on Unit Ball Graphs in \mathbb{R}^d . Our results in this section are based on an observation that if for some graph class \mathcal{G} the size of an optimum solution for a problem Π and the number of vertices in the input graph are linearly related then to obtain EPTAS it is sufficient that \mathcal{G} has sublinear treewidth rather than truly sublinear treewidth. The crux of this result is based on the following adaptation of the decomposition lemma proved in [26, Lemma 3.2].

Lemma 11. *Let \mathcal{G} be a hereditary graph class of sublinear treewidth with parameter $\lambda < 1$, that is, for every $G \in \mathcal{G}$, $\text{tw}(G) = O(|V(G)|^\lambda)$. For every $\epsilon < 1$ there is γ such that for any $G \in \mathcal{G}$ there is $X \subseteq V(G)$ satisfying $|X| \leq \epsilon|V(G)|$ and for every connected component C of $G \setminus X$, we have that $|C| \leq \gamma$. Moreover X can be computed from G in polynomial time.*

Proof. Let the number of vertices of G be n , that is, $|V(G)| = n$. For any $\gamma \geq 1$, define $T_\gamma : \mathbb{N} \rightarrow \mathbb{N}$ such that $T_\gamma(n)$ is the smallest integer such that if $G \in \mathcal{G}$ and $|V(G)| \leq n$, then there is a $X \subseteq V(G)$ of size at most $T_\gamma(n)$ such that for every connected component C of $G \setminus X$ we have $|C| \leq \gamma$. Informally, $T_\gamma(n)$ is the minimum size of a vertex set X such that every connected component C of $G \setminus X$ has at most γ vertices. Furthermore, since \mathcal{G} is a hereditary graph class of sublinear treewidth with parameter λ there exists a constant β such that $\text{tw}(G) \leq \beta n^\lambda$. We will make choices for the constants δ and γ and ρ based on λ , β and ϵ . Our aim is to show that $T_\gamma(n) \leq \epsilon n$ for every n .

Observe that for any numbers $a > 0$, $b > 0$, we have $a^\lambda + b^\lambda > (a + b)^\lambda$ since $\lambda < 1$. Thus we have $\rho = \min_{1/3 \leq \alpha \leq 2/3} \alpha^\lambda + (1 - \alpha)^\lambda > 1$. We choose $\delta = \frac{(2\epsilon+1)(\beta+1)}{\rho-1}$ and $\gamma = \left(\frac{3\delta}{\epsilon}\right)^{\frac{1}{1-\lambda}}$. If $|V(G)| \leq \gamma$ then we set $X = \emptyset$, so $T_\gamma(n) = 0 \leq \epsilon n$ for $n \leq \gamma$. We now show that if $n \geq \gamma/3$ then $T_\gamma(n) = 0 \leq \epsilon n - \delta n^\lambda$ by induction on n . For the base case if $\gamma/3 \leq n \leq \gamma$ then the choice of γ implies the following inequality.

$$\epsilon n - \delta n^\lambda \geq \epsilon \frac{\gamma}{3} - \delta \gamma^\lambda \geq 0 = T_\gamma(n)$$

We now consider $T_\gamma(n)$ for $n > \gamma$. We know that the treewidth of G is at most βn^λ . Construct a weight function $w : V(G) \rightarrow \mathbb{N}$ such that $w(v) = 1$, for all $v \in V(G)$. By Lemma 9, there is a partition of $V(G)$ into L, S and R such that $|S| \leq \beta n^\lambda + 1$, $N(L) \subseteq S$, $N(R) \subseteq S$, $|L| \leq 2n/3$ and $|R| \leq 2n/3$. Deleting S from the graph G yields two graphs $G[L]$ and $G[R]$ with no edges between them. Thus we put S into X and then proceed recursively in $G[L]$ and $G[R]$. This yields the following recurrence for T_γ .

$$T_\gamma(n) \leq \max_{1/3 \leq \alpha \leq 2/3} T(\alpha n + \beta n^\lambda + 1) + T((1 - \alpha)n + \beta n^\lambda + 1) + \beta n^\lambda + 1.$$

Observe that since $n \geq \gamma$ we have $\alpha n \geq \gamma/3$ and $(1 - \alpha)n \geq \gamma/3$. The induction hypothesis then yields

the following inequality.

$$\begin{aligned}
T_\gamma(n) &\leq \max_{1/3 \leq \alpha \leq 2/3} T(\alpha n + \beta n^\lambda + 1) + T((1 - \alpha)n + \beta n^\lambda + 1) + \beta n^\lambda + 1 \\
&\leq \max_{1/3 \leq \alpha \leq 2/3} \epsilon n - \delta(\alpha n)^\lambda - \delta((1 - \alpha)n)^\lambda + (2\epsilon + 1)(\beta n^\lambda + 1) \\
&\leq \max_{1/3 \leq \alpha \leq 2/3} \epsilon n - \delta n^\lambda (\alpha^\lambda + (1 - \alpha)^\lambda) + (2\epsilon + 1)(\beta n^\lambda + 1) \\
&\leq \epsilon n - \delta n^\lambda - \delta(\rho - 1)n^\lambda + (2\epsilon + 1)(\beta n^\lambda + 1) \\
&\leq \epsilon n - \delta n^\lambda.
\end{aligned}$$

The last inequality holds whenever $\delta(\rho - 1)n^\lambda \geq (2\epsilon + 1)(\beta n^\lambda + 1)$, which is ensured by the choice of δ and the fact that $n^\lambda \geq 1$. Thus $T_\gamma(n) \leq \epsilon n$ for all n . Hence there exists a set X of size at most ϵn such that for every component C of $G \setminus X$ we have $|C| \leq \gamma$.

What remains is to show that X can be computed from G in polynomial time. The inductive proof can be converted into a recursive algorithm. The only computationally hard step of the proof is the construction of a tree-decomposition of G in each inductive step. Instead of computing the treewidth exactly we use the $d^* \sqrt{\log \text{tw}(G)}$ -approximation algorithm by Feige et al. [25], where d^* is a fixed constant. Thus when we partition $V(G)$ into L , S , and R using Lemma 9, the upper bound on the size of S will be $d^*(\beta n^\lambda + 1)\sqrt{\log(\beta n^\lambda)}$ instead of $\beta n^\lambda + 1$. However, for any $\lambda < \lambda' < 1$ there is a β' such that $d^*(\beta n^\lambda)\sqrt{\log(\beta n^\lambda)} < \beta' n^{\lambda'}$. Now we can apply the above analysis with β' instead of β and λ' instead of λ to bound the size of the set X output by the algorithm. This concludes the proof of the lemma. \square

Using Lemma 11 we can obtain the following analogue of Proposition 2 ([26, Theorem 4.1]).

Theorem 8. *Let Π be a reducible graph optimization problem and let \mathcal{G} be a class of graphs with sublinear treewidth such that for every $G \in \mathcal{G}$, $\pi(G) = \Omega(|V(G)|)$. Then Π has an EPTAS on \mathcal{G} .*

Proof. Let G be the input to Π , $|V(G)| = n$ and $\epsilon > 0$ be fixed. Since for every $G \in \mathcal{G}$, $\pi(G) = \Omega(n)$, we have that $\pi(G) \geq \rho_1 n$, for a fixed constant ρ_1 . Furthermore, since \mathcal{G} is a hereditary graph class of sublinear treewidth with parameter λ , there exists a constant β such that $\text{tw}(G) \leq \beta n^\lambda$. Let ϵ' be a constant to be selected later. By Lemma 11, there exist γ , $\lambda' < 1$ and β' depending on ϵ' , λ and β such that given G a set X with the following properties can be found in polynomial time. First $|X| \leq \epsilon' n$, and secondly for every component C of $G \setminus X$ we have that $|C| \leq \gamma$. Thus $\text{tw}(G \setminus X) = \tau \leq \beta' \gamma^{\lambda'}$. Since Π is reducible, there exists a MIN/MAX-CMSO problem Π' , a constant ρ_2 and a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that:

1. there is a polynomial time algorithm that given G and $X \subseteq V(G)$ outputs G' such that $|\pi'(G') - \pi(G)| \leq \rho_2 |X|$ and $\text{tw}(G') \leq f(\tau)$,
2. there is a polynomial time algorithm that given G and $X \subseteq V(G)$, G' and a vertex (edge) set S' such that $P_{\Pi'}(G', S')$ holds outputs S such that $\phi_\Pi(G, S)$ holds and $|\kappa_\Pi(G, S) - |S'|| \leq \rho_2 |X|$.

We construct G' from G and X using the first polynomial time algorithm. Since $\text{tw}(G') \leq f(\tau)$ we can use an extended version of Courcelle's theorem [14, 15] given by Borie et al. [8] to find an optimal solution S' to Π' in $g(\epsilon')|V(G')|$ time. By the properties of the first polynomial time algorithm, $||S'| - \pi(G)| \leq \rho |X|$ where $\rho = \max(\rho_1, \rho_2)$. We now use the second polynomial time algorithm to construct a solution S to Π from G , X , G' and S' . The properties of the second algorithm ensure $\phi_\Pi(G, S)$ holds and that $|\kappa_\Pi(G, S) - |S'|| \leq \rho |X|$, and hence $|\kappa_\Pi(G, S) - \pi(G)| \leq 2\rho |X| \leq 2\rho^2 \epsilon' \pi(G)$. Choosing $\epsilon' = \frac{\epsilon}{2\rho^2}$ yields $|\kappa_\Pi(G, S) - \pi(G)| \leq \epsilon \pi(G)$, proving the theorem. \square

Lemma 12. *Let G be an intersection graph of unit balls in \mathbb{R}^d , for a fixed d . If G does not contain an isolated vertex then the size of minimum (connected) vertex cover is at least $|V(G)|/f(d)$, where $f(d) = 2(2^{0.401d(1+o(1))} + 1)$.*

Proof. Let G be an intersection graph of unit balls in \mathbb{R}^d , for a fixed d . For our proof we need the concept of *kissing number*. The kissing number τ_d is the maximum number of non overlapping d -dimensional unit balls of equal size that can touch a unit ball in \mathbb{R}^d . Kabatiansky and Levenshtein [31] showed that $\tau_d \leq 2^{0.401d(1+o(1))}$. This implies that for any vertex $v \in V(G)$, $N(v)$ does not contain an independent set of size bigger than $\tau_d + 1$.

Given a graph G we compute a maximal matching, say M . Clearly the size of M is a lower bound on the size of a minimum (connected) vertex cover. Let $V(M)$ be the set of end points of edges in M and $I = V(G) \setminus V(M)$. Clearly I is an independent set. Furthermore every vertex in I is adjacent to some vertex in $V(M)$. Hence we have that $|I| \leq |V(M)|(\tau_d + 1)$. This implies that $|V(G)| = |V(M)| + |I| \leq 2|M| + 2|M|(\tau_d + 1)$. The last inequality implies the lemma. \square

Finally, we note that every graph G , that is, an intersection graph of unit balls in \mathbb{R}^d , with maximum clique size Δ has the property that every point in \mathbb{R}^d is in at most Δ unit balls. This together with result from [35] implies that the treewidth of G is $c_d \Delta^{1/d} |V(G)|^{1-\frac{1}{d}}$, where c_d is a constant depending only on d . This implies that an intersection graph of unit balls in \mathbb{R}^d with bounded maximum clique has sublinear treewidth. So an EPTAS for CONNECTED VERTEX COVER and VERTEX COVER can be obtained along the similar lines as in Theorems 5 and 6 and finally uses Theorem 8 instead of Theorem 2 to arrive to the following result.

Theorem 9. CONNECTED VERTEX COVER and VERTEX COVER admit an EPTAS on unit ball graphs of fixed dimension.

We can also obtain EPTASs for CONNECTED VERTEX COVER and VERTEX COVER on disk graphs, as on disk graphs of bounded clique size we have that the size of an optimum solution and the number of vertices in the input graph are linearly related.

Parameterized Subexponential Time Algorithms. In this section we show how to obtain parameterized subexponential time algorithm for several problems. Formally, a *parameterization* of a problem is assigning an integer k to each input instance and a parameterized problem is *fixed-parameter tractable* (FPT) if there is an algorithm that solves the problem in time $f(k) \cdot |I|^{O(1)}$, where $|I|$ is the size of the input and f is an arbitrary computable function. We say that a parameterized problem has a parameterized subexponential algorithm if it is solvable in time $2^{o(k)} \cdot |I|^{O(1)}$.

Our basic idea is to find a “large” clique and guess the intersection of an optimal solution with this clique. We recursively do this until we do not have a large clique. Once we do not have large clique we have that the maximum degree of the graph is bounded and thus the treewidth comes into picture. At that point we use dynamic programming on graphs of bounded treewidth to solve the problem optimally. We exemplify our approach on p -FEEDBACK VERTEX SET. In this problem we are given a graph G and a positive integer k and the question is to check whether there is a subset $F \subseteq V(G)$, $|F| \leq k$, such that $G \setminus F$ is acyclic. The set F is called feedback vertex set of G .

Theorem 10. p -FEEDBACK VERTEX SET admits a parameterized subexponential time algorithm on unit disk graphs and map graphs.

Proof. We give a subexponential time parameterized algorithm on map graphs. An algorithm on unit disk graphs is similar. The algorithm proceeds as follows. Given an instance (G, k) , it finds a maximum clique C of G . Recall that we can find a maximum clique in unit disk graphs and map graphs in polynomial time [11, 12, 13, 37]. If $|C| > k + 2$, then we return that G does not have feedback vertex set of size at most k . Next we check whether $|C| \leq k^\epsilon$ (ϵ to be fixed later). If yes then by Theorem 2 we know that $\text{tw}(G) \leq O(k^{0.5+\epsilon})$. In this case we apply the known algorithm for FEEDBACK VERTEX SET, that given a tree decomposition of width t of a graph G on n vertices, finds a minimum sized feedback vertex set in time $2^{O(t \log t)} n^{O(1)}$. Hence in this case the running time of our algorithm will be $2^{O(k^{0.5+\epsilon} \log k)} n^{O(1)}$. Now we consider the case when $|C| > k^\epsilon$. We know that for any feedback

vertex set F of G , we have that $|F \cap C| \geq |C| - 2$. So we guess the intersection $X = F \cap C$ and recursively solve the problem on $(G \setminus X, k - |X|)$. If for any guess we have an yes answer we return yes, else, we return no. The running time of this step is guided by the following recurrence $T(k) \leq \binom{|C|}{2} \cdot T(k - (|C| - 2)) + |C| \cdot T(k - (|C| - 1)) + T(k - |C|)$, where the terms $\binom{|C|}{2} \cdot T(k - (|C| - 2))$, $|C| \cdot T(k - (|C| - 1))$, $T(k - |C|)$ correspond to choosing $|C| - 2$ vertices in F from C , $|C| - 1$ vertices in F from C and $|C|$ vertices in F from C , respectively. Roughly, $T(k) \leq 3|C|^2 \cdot T(k - |C|)$. This asymptotically solves to $(3|C|)^{2k/|C|}$, which is equal to $2^{O(\frac{2k \log |C|}{|C|})} \leq 2^{O(\frac{2k \log k}{|C|})}$. Hence as $|C|$ increases we have that the function $2^{O(\frac{2k \log k}{|C|})}$ decreases. Thus the worst case running time is achieved when $|C| = k^\epsilon$ and hence this is equal to $2^{O(k^{1-\epsilon} \log k)}$. Now we choose ϵ in a way that the running time for branching on clique is same as when we run a dynamic programming algorithm on graphs of bounded treewidth. Thus we choose an ϵ such that $2^{O(k^{1-\epsilon} \log k)} = 2^{O(k^{0.5+\epsilon})}$. This gives us that $\epsilon = 1/4$ is asymptotically best possible. Thus our algorithm runs in time $2^{O(k^{0.75} \log k)} n^{O(1)} = 2^{o(k)} n^{O(1)}$. \square

Next we show that in fact p -(CONNECTED) VERTEX COVER admits a parameterized subexponential time algorithms on Unit Ball Graphs in \mathbb{R}^d .

Theorem 11. p -CONNECTED VERTEX COVER and p -VERTEX COVER admit a parameterized subexponential time algorithm on unit ball graphs of fixed dimension.

Proof. Our algorithm for p -(CONNECTED) VERTEX COVER follows along the same line as for p -FEEDBACK VERTEX SET. We only outline an algorithm for p -CONNECTED VERTEX COVER here.

The algorithm proceeds as follows. Given an instance (G, k) , we first check whether $k \geq |V(G)|/f(d)$, where $f(d) = 2(2^{0.401d(1+o(1))} + 1)$. By Lemma 12 we know that if $k < |V(G)|/f(d)$, then there is no connected vertex cover of size at most k and hence the answer is no. Else we have that $|V(G)| = O(k)$. To implement our algorithm we need a slight generalization of problem. We keep a triple (G', k, X) for this problem, where G' is the current graph and the objective is to find a set $F \subseteq V(G')$ such that $|F| \leq k$, F is a vertex cover of G' and $G[X \cup F]$ is a connected vertex cover of G . Essentially the graph G' will be obtained after branching on cliques and the set X will store the partially constructed solution so far. This allows us to check connectedness in the whole graph G . Now the algorithm finds a maximum clique C of G' . If $|C| > k + 1$, then we return that G' does not have a desired set F of size at most k . Next we check whether $|C| \leq k^\epsilon$ (ϵ to be fixed later).

We first consider the case when $|C| > k^\epsilon$. We know that for any vertex cover F of G' , we have that $|F \cap C| \geq |C| - 1$. So we guess the intersection $Z = F \cap C$ and recursively solve the problem on $(G' \setminus Z, k - |Z|, X \cup Z)$. If for any guess we have an yes answer we return yes else we return no. The running time of this step is guided by the following recurrence $T(k) \leq |C| \cdot T(k - (|C| - 1)) + T(k - |C|)$, where the terms $|C| \cdot T(k - (|C| - 1))$, $T(k - |C|)$ correspond to choosing $|C| - 1$ vertices in F from C and $|C|$ vertices in F from C , respectively. Roughly $T(k) \leq (2|C|)T(k - |C|)$. This asymptotically solves to $(2|C|)^{k/|C|}$, which is equal to $2^{O(\frac{k \log |C|}{|C|})} \leq 2^{O(\frac{k \log k}{|C|})}$. Hence as $|C|$ increases we have that the function $2^{O(\frac{k \log k}{|C|})}$ decreases. Thus the worst case running time is achieved when $|C| = k^\epsilon$ and hence this is equal to $2^{O(k^{1-\epsilon} \log k)}$.

In the other case we have that $|C| \leq k^\epsilon$. As discussed before Lemma 9, by using result from [35] we have that the treewidth of G' is $c_d k^{\epsilon/d} |V(G)|^{1-\frac{1}{d}} = O(k^{1-(1-\epsilon)\frac{1}{d}})$, where c_d is a constant depending only on d . In this case we apply a modification of known algorithm for CONNECTED VERTEX COVER, that given a tree decomposition of width t of a graph G^* on n vertices, finds a minimum sized connected vertex cover in time $2^{O(t \log t)} n^{O(1)}$ [36]. To solve our problem we do as follows. We first upper bound the number of connected components, η_X , in $G[X]$ by $k^{1-\epsilon}$. Recall that X has been constructed by branching on cliques of size at least $k^\epsilon + 1$ and thus from each such clique we have at least k^ϵ vertices in X and vertices from one clique are in one component. Thus $\eta_X \leq k/k^\epsilon = k^{1-\epsilon}$. Now we construct a graph G^* as follows. Consider the graph $G[X \cup V(G')]$ and contract every connected component in $G[X]$ to a single vertex. Now in the graph G^* the objective is to find a connected vertex cover of size

at most $k + \eta_X$ such that it contains all the vertices corresponding to connected components in $G[X]$. Now the $\text{tw}(G^*) \leq \text{tw}(G') + \eta_X \leq O(k^{1-(1-\epsilon)\frac{1}{d}} + k^{1-\epsilon})$. Hence in this case the running time of our algorithm is $2^{O((k^{1-(1-\epsilon)\frac{1}{d}} + k^{1-\epsilon}) \log k)} n^{O(1)}$.

Now we choose ϵ in a way that the running time for branching on clique is same as when we run a dynamic programming algorithm on graphs of bounded treewidth. Thus, we choose an ϵ such that $2^{O(k^{1-\epsilon} \log k)} = 2^{O((k^{1-(1-\epsilon)\frac{1}{d}} + k^{1-\epsilon}) \log k)}$. This gives us that $\epsilon = 1/(d + 1)$ is asymptotically best possible. Thus, our algorithm runs in time $2^{O((k^{1-(1-\epsilon)\frac{1}{d}} + k^{1-\epsilon}) \log k)} n^{O(1)} = 2^{o(k)} n^{O(1)}$ for every fixed d . This gives us the desired result. \square

Tractability Borders. It is natural to ask how far our approach can be generalized, and in particular, whether many of the problems discussed so far have EPTASs and parameterized subexponential time algorithms on unit ball graphs in dimension higher than two. In this section we show that one should not expect equally general results for unit ball graphs of dimension at least three. In particular, we show that FEEDBACK VERTEX SET on Unit Ball Graphs in \mathbb{R}^3 does not have an EPTAS unless $P = NP$, and that the problem does not admit a subexponential time parameterized algorithm under the Exponential Time Hypothesis of Impagliazzo, Paturi and Zane [30].

Theorem 12. FEEDBACK VERTEX SET on unit ball graphs in \mathbb{R}^3 does not admit a PTAS unless $P = NP$, and has no subexponential time parameterized algorithm unless the Exponential Time Hypothesis fails.

A unit ball model of H in \mathbb{R}^d is a map $f : V(H) \rightarrow \mathbb{R}^d$ such that u and v are adjacent iff the euclidean distance between $f(u)$ and $f(v)$ is at most 1. In the construction it is much more convenient to work with this alternate definition of unit ball graphs rather than saying that $f(u)$ and $f(v)$ is at most 2 and hence we use this alternate definition in this section. In our constructions no two vertices will map to the same point, and thus we will often refer to vertices in H by the points in \mathbb{R}^d which they map to. For the proof of Theorem 12 we need the following lemmas. It appears that the following lemma can easily be derived from the results in [23] about the three dimensional orthogonal graph drawings. However, since we could not find this result explicitly, we give a proof here for completeness.

Lemma 13. For any graph G on n vertices of maximum degree 6, there is a unit ball graph H on $O(n^2)$ vertices such that H is a subdivision of G . Furthermore, H and a unit ball model of H in \mathbb{R}^3 can be constructed from G in polynomial time.

The proof of Lemma 13 is straightforward, but somewhat tedious.

Proof. In this construction we envision the x -axis as being horizontal with positive direction towards the right, the z -axis being vertical with positive direction upwards. The intuition behind the proof is that every vertex of G is assigned its own “fat” x - z plane. The edges of G are routed parallel to the y axis in the $y - x$ plane with $z = 0$, and in each “fat” x - z plane we ensure that the edges connect to their corresponding vertex. This local routing of edge endpoints to a vertex happens above the $y - x$ plane with $z = 0$ and does not interfere with the global routing of the edges.

For a point with integer coordinates (x, y, z) and integer ℓ define the set $L[x, y, z]_x^\ell$ to be $\{(x + x', y, z) : |x'| + |\ell - x'| = |\ell|\}$. In particular, if ℓ is positive then $L[x, y, z]_x^\ell$ contains $\{(x, y, z), (x + 1, y, z), (x + 2, y, z), \dots, (x + \ell, y, z)\}$, while if ℓ is negative then $L[x, y, z]_x^\ell$ contains $\{(x, y, z), (x - 1, y, z), (x - 2, y, z), \dots, (x - \ell, y, z)\}$. Similarly we define $L[x, y, z]_y^\ell$ to be $\{(x, y + y', z) : |y'| + |\ell - y'| = |\ell|\}$ and $L[x, y, z]_z^\ell$ to be $\{(x, y, z + z') : |z'| + |\ell - z'| = |\ell|\}$. Given three integers x, y, z , the graph

$P[x, y, z]$ corresponds to the point set

$$\begin{aligned}
P[x, y, z] &= L[x, y, z]_{\mathbf{z}}^{-2} \cup L[x, y, z - 2]_{\mathbf{x}}^2 \\
&\cup L[x, y, z]_{\mathbf{x}}^2 \\
&\cup L[x, y, z]_{\mathbf{z}}^2 \cup L[x, y, z + 2]_{\mathbf{x}}^2 \\
&\cup L[x, y, z]_{\mathbf{x}}^{-2} \cup L[x - 2, y, z]_{\mathbf{z}}^4 \cup L[x - 2, y, z + 4]_{\mathbf{x}}^4 \\
&\cup L[x, y, z]_{\mathbf{y}}^{-2} \cup L[x, y - 2, z]_{\mathbf{z}}^6 \cup L[x, y - 2, z + 6]_{\mathbf{y}}^2 \cup L[x, y, z + 6]_{\mathbf{x}}^2 \\
&\cup L[x, y, z]_{\mathbf{y}}^2 \cup L[x, y + 2, z]_{\mathbf{z}}^8 \cup L[x, y + 2, z + 8]_{\mathbf{y}}^{-2} \cup L[x, y, z + 8]_{\mathbf{x}}^2
\end{aligned}$$

The set $P[x, y, z]$ corresponds to a vertex of degree 6 in $[x, y, z]$, and there are 6 paths, each starting in (x, y, z) and ending in $(x + 2, y, z - 2)$, $(x + 2, y, z)$, $(x + 2, y, z + 2)$, $(x + 2, y, z + 4)$, $(x + 2, y, z + 6)$ and $(x + 2, y, z + 8)$ respectively. The y -coordinate of any intermediate point on the paths is always between $y - 2$ and $y + 2$. Any points that are generated twice still correspond only to one single vertex.

For an integer y and six integers $x_1 < x_2 < \dots < x_6$ such that $x_{i+1} - x_i \geq 2$ we define $\mathbf{P}[y, x_1, x_2, x_3, x_4, x_5, x_6]$ to be the point set

$$\begin{aligned}
\mathbf{P}[y, x_1, x_2, x_3, x_4, x_5, x_6] &= P[-2, y, 12] \\
&\cup \bigcup_{i=1}^6 L[0, y, 10 + 2(i - 1)]_{\mathbf{x}}^{x_i} \cup L[x_i, y, 10 + 2(i - 1)]_{\mathbf{z}}^{-10 - 2(i - 1)}
\end{aligned}$$

The set $\mathbf{P}[y, x_1, x_2, x_3, x_4, x_5, x_6]$ corresponds to a vertex of degree 6 in $[-2, y, 12]$ with 6 paths starting in this vertex and ending in $[x_i, y, 0]$ for $1 \leq i \leq 6$. The y -coordinate of the intermediate vertices on the path is between $y - 2$ and $y + 2$. In this sense, the paths corresponding to the vertex in $[-2, y, 12]$ are routed in a “fat” $x - z$ -plane.

We are now ready to construct H given G . We give the construction for 6-regular graphs G and then explain how to modify the construction to the case when G has maximum degree 6. We label the vertices in G by v_1, \dots, v_n and the edges of G by e_1, \dots, e_m with $m \leq 3n$. For every $i \leq m$ define $a(i)$ and $b(i)$ such that the endpoints of the edge e_i are $v_{a(i)}$ and $v_{b(i)}$ respectively. Now, for every vertex v_i let $x_1^i < x_2^i < \dots < x_6^i$ be integers so that v_i is incident to the edges $e_{x_j^i}$ for $1 \leq j \leq 6$. For every vertex v_i we add the point set $\mathbf{P}[10i, 2x_1^i, 2x_2^i, 2x_3^i, 2x_4^i, 2x_5^i, 2x_6^i]$. Finally for every edge e_i we add the set $L[2i, 10a(i), 0]_{\mathbf{y}}^{10(b(i) - a(i))}$. This concludes the construction of H .

It is easy to see that H can be constructed from G in polynomial time. Furthermore, it is easy to verify that H has $O(n^2)$ vertices since $m \leq 3n$. To see that H is a subdivision of G observe that when G has an edge e_t between v_i and v_j , in H there is a path from the point $[-2, 10i, 12]$ through $[2t, 10i, 0]$ and $[2t, 10j, 0]$ to the point $[-2, 10j, 12]$. This concludes the proof of the lemma. \square

Lemma 14. *There is a polynomial time algorithm that given a graph G on n vertices of maximum degree 3 outputs a unit ball graph H together with a unit ball model of H in \mathbb{R}^3 , such that given any vertex cover C of G , a feedback vertex set S of H of size at most $|C|$ can be computed in polynomial time, and given any feedback vertex set S of H , a vertex cover C of G of size at most $|S|$ can be computed in polynomial time.*

Proof. Given G we start by applying the well-known construction for transforming instances of VERTEX COVER to instances of FEEDBACK VERTEX SET. We construct G' from G by adding a vertex x_{uv} for every edge uv of G and making x_{uv} adjacent to u and to v . Since the maximum degree of G was 3, the maximum degree of G' is 6. Now we apply Lemma 13 to G' obtain the graph H and a unit ball model of H . Every vertex cover C of G is a feedback vertex set of G' , and since H is a subdivision of G' , every vertex cover of G is a feedback vertex set of H . For the reverse direction, it is well-known that given a feedback vertex set S in a graph, one can find in polynomial time a feedback vertex set S' of size at most $|S|$ such that all vertices in S' have degree at least 3 [6]. Let S' be a feedback vertex set of H such that

every vertex in S' has degree at least 3 in H . Then, every vertex in S' is also a vertex in G . We claim that S' is a vertex cover of G . Let uv be an edge in G . Then u, x_{uv}, v is a cycle in G' and since H is a subdivision of G' , H contains a cycle going through u, x_{uv} and v where all vertices in the cycle except u and v have degree at most 2. Since S' is a feedback vertex set of H' containing no vertices of degree less than 3, S' contains either u or v . Hence S' is a vertex cover of G . \square

If a subexponential time parameterized algorithm for FEEDBACK VERTEX SET on unit ball graphs in \mathbb{R}^3 existed we could combine it with Lemma 14 to get a subexponential time algorithm for VERTEX COVER on graphs of maximum degree 3. Similarly, a PTAS for FEEDBACK VERTEX SET on unit ball graphs in \mathbb{R}^3 could be combined with Lemma 14 to yield a PTAS for VERTEX COVER on graphs of maximum degree 3. Since VERTEX COVER is known not to admit a $(1 + \epsilon)$ -factor approximation algorithm, for some fixed $\epsilon > 0$, on graphs of degree at most 3 unless $P = NP$ [3], and not to have subexponential time parameterized algorithms on graphs of degree at most 3 under the Exponential Time Hypothesis [30], we obtain Theorem 12.

References

- [1] P. K. AGARWAL, M. VAN KREVELD, AND S. SURI, *Label placement by maximum independent set in rectangles*, Computational Geometry, 11 (1998), pp. 209–218.
- [2] J. ALBER AND J. FIALA, *Geometric separation and exact solutions for the parameterized independent set problem on disk graphs*, J. Algorithms, 52 (2004), pp. 134–151.
- [3] P. ALIMONTI AND V. KANN, *Some APX-completeness results for cubic graphs*, Theor. Comput. Sci., 237 (2000), pp. 123–134.
- [4] S. ARNBORG, J. LAGERGREN, AND D. SEESE, *Easy problems for tree-decomposable graphs*, J. Algorithms, 12 (1991), pp. 308–340.
- [5] B. S. BAKER, *Approximation algorithms for NP-complete problems on planar graphs*, J. ACM, 41 (1994), pp. 153–180.
- [6] R. BAR-YEHUDA, D. GEIGER, J. NAOR, AND R. M. ROTH, *Approximation algorithms for the feedback vertex set problem with applications to constraint satisfaction and Bayesian inference*, SIAM Journal on Computing, 27 (1998), pp. 942–959.
- [7] H. L. BODLAENDER, *A partial k -arboretum of graphs with bounded treewidth*, Theor. Comp. Sc., 209 (1998), pp. 1–45.
- [8] R. B. BORIE, R. G. PARKER, AND C. A. TOVEY, *Automatic generation of linear-time algorithms from predicate calculus descriptions of problems on recursively constructed graph families*, Algorithmica, 7 (1992), pp. 555–581.
- [9] T. M. CHAN, *Polynomial-time approximation schemes for packing and piercing fat objects*, J. Algorithms, 46 (2003), pp. 178–189.
- [10] Z.-Z. CHEN, *Approximation algorithms for independent sets in map graphs*, J. Algorithms, 41 (2001), pp. 20–40.
- [11] Z.-Z. CHEN, E. GRIGNI, AND C. H. PAPADIMITRIOU, *Planar map graphs*, in Proceedings of the 30th Annual ACM Symposium on the Theory of Computing (STOC '98), ACM, 1998, pp. 514–523.
- [12] ———, *Map graphs*, Journal of the Association for Computing Machinery, 49 (2002), pp. 127–138.
- [13] B. N. CLARK, C. J. COLBOURN, AND D. S. JOHNSON, *Unit disk graphs*, Discrete Mathematics, 86 (1990), pp. 165–177.
- [14] B. COURCELLE, *The monadic second-order logic of graphs I: Recognizable sets of finite graphs*, Information and Computation, 85 (1990), pp. 12–75.
- [15] ———, *The expression of graph properties and graph transformations in monadic second-order logic*, Handbook of Graph Grammars, (1997), pp. 313–400.
- [16] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Fixed-parameter algorithms for (k, r) -center in planar graphs and map graphs*, ACM Trans. Algorithms, 1 (2005), pp. 33–47.
- [17] E. D. DEMAINE, F. V. FOMIN, M. HAJIAGHAYI, AND D. M. THILIKOS, *Subexponential parameterized algorithms on bounded-genus graphs and H -minor-free graphs*, J. ACM, 52 (2005), pp. 866–893.

- [18] E. D. DEMAINE AND M. HAJIAGHAYI, *Bidimensionality: New connections between FPT algorithms and PTASs*, in Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2005), ACM-SIAM, 2005, pp. 590–601.
- [19] E. D. DEMAINE AND M. HAJIAGHAYI, *Linearity of grid minors in treewidth with applications through bidimensionality*, *Combinatorica*, 28 (2008), pp. 19–36.
- [20] E. D. DEMAINE, M. HAJIAGHAYI, AND K. KAWARABAYASHI, *Algorithmic graph minor theory: Improved grid minor bounds and Wagner’s contraction*, *Algorithmica*, 54 (2009), pp. 142–180.
- [21] R. DIESTEL, *Graph Theory*, Springer-Verlag, Heidelberg, third ed., 2005.
- [22] A. DUMITRESCU AND J. PACH, *Minimum clique partition in unit disk graphs*, *Graph. Comb.*, 27 (2011), pp. 399–411.
- [23] P. EADES, A. SYMVONIS, AND S. WHITESIDES, *Three-dimensional orthogonal graph drawing algorithms*, *Discrete Applied Mathematics*, 103 (2000), pp. 55–87.
- [24] T. ERLEBACH, K. JANSEN, AND E. SEIDEL, *Polynomial-time approximation schemes for geometric intersection graphs*, *SIAM J. Comput.*, 34 (2005), pp. 1302–1323.
- [25] U. FEIGE, M. HAJIAGHAYI, AND J. R. LEE, *Improved approximation algorithms for minimum-weight vertex separators*, in Proceedings of the 37th annual ACM Symposium on Theory of computing (STOC 2005), New York, 2005, ACM Press, pp. 563–572.
- [26] F. V. FOMIN, D. LOKSHTANOV, V. RAMAN, AND S. SAURABH, *Bidimensionality and EPTAS*, in Proceedings of the 22nd annual ACM-SIAM symposium on Discrete algorithms (SODA 2011), SIAM, 2011, pp. 748–759.
- [27] F. V. FOMIN, D. LOKSHTANOV, S. SAURABH, AND D. M. THILIKOS, *Bidimensionality and kernels*, in Proceedings of the 21st annual ACM-SIAM symposium on Discrete algorithms (SODA 2010), ACM-SIAM, 2010, pp. 503–510.
- [28] D. S. HOCHBAUM AND W. MAASS, *Approximation schemes for covering and packing problems in image processing and vlsi*, *J. ACM*, 32 (1985), pp. 130–136.
- [29] H. B. HUNT, III, M. V. MARATHE, V. RADHAKRISHNAN, S. S. RAVI, D. J. ROSENKRANTZ, AND R. E. STEARNS, *NC-approximation schemes for NP- and PSPACE-hard problems for geometric graphs*, *J. Algorithms*, 26 (1998), pp. 238–274.
- [30] R. IMPAGLIAZZO, R. PATURI, AND F. ZANE, *Which problems have strongly exponential complexity*, *Journal of Computer and System Sciences*, 63 (2001), pp. 512–530.
- [31] G. KABATIANSKY AND V. LEVENSHTAIN, *Bounds for packings on a sphere and in space*, *Problemy Peredachi Informatsii*, 14 (1978), pp. 3–25.
- [32] F. KUHN, R. WATTENHOFER, AND A. ZOLLINGER, *Ad hoc networks beyond unit disk graphs*, *Wireless Networks*, 14 (2008), pp. 715–729.
- [33] M. V. MARATHE, H. BREU, H. B. HUNT, III, S. S. RAVI, AND D. J. ROSENKRANTZ, *Simple heuristics for unit disk graphs*, *Networks*, 25 (1995), pp. 59–68.
- [34] D. MARX, *Parameterized complexity and approximation algorithms*, *The Computer Journal*, 51 (2008), pp. 60–78.
- [35] G. L. MILLER, S.-H. TENG, W. P. THURSTON, AND S. A. VAVASIS, *Separators for sphere-packings and nearest neighbor graphs*, *J. ACM*, 44 (1997), pp. 1–29.

- [36] H. MOSER, *Exact algorithms for generalizations of vertex cover*, Master's thesis, Institut für Informatik, Friedrich-Schiller-Universität, 2005.
- [37] V. RAGHAVAN AND J. SPINRAD, *Robust algorithms for restricted domains*, J. Algorithms, 48 (2003), pp. 160–172.
- [38] N. ROBERTSON, P. D. SEYMOUR, AND R. THOMAS, *Quickly excluding a planar graph*, J. Comb. Theory Series B, 62 (1994), pp. 323–348.
- [39] P. D. SEYMOUR AND R. THOMAS, *Graph searching and a minimax theorem for tree-width*, J. Comb. Theory Series B, 58 (1993), pp. 239–257.
- [40] M. THORUP, *Map graphs in polynomial time*, in Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS 1998), IEEE Computer Society, 1998, pp. 396–405.
- [41] E. J. VAN LEEUWEN, *Optimization and approximation on systems of geometric objects*, PhD thesis, De Vries Instituut, Netherlands, 2009.
- [42] J. XU AND B. BERGER, *Fast and accurate algorithms for protein side-chain packing*, J. ACM, 53 (2006), pp. 533–557.