

Approximation Algorithms for Polynomial-Expansion and Low-Density Graphs*

Sariel Har-Peled[†]

Kent Quanrud[‡]

August 14, 2017

Abstract

We investigate the family of intersection graphs of low density objects in low dimensional Euclidean space. This family is quite general, includes planar graphs, and in particular is a subset of the family of graphs that have polynomial expansion.

We present efficient $(1 + \varepsilon)$ -approximation algorithms for polynomial expansion graphs, for unweighted **Independent Set**, **Set Cover**, and **Dominating Set** problems, among others, and these results seem to be new. Naturally, PTAS's for these problems are known for subclasses of this graph family. These results have immediate applications in the geometric domain. For example, the new algorithms yield the only PTAS known for covering points by fat triangles (that are shallow).

We also prove corresponding hardness of approximation for some of these optimization problems, characterizing their intractability with respect to density. For example, we show that there is no PTAS for covering points by fat triangles if they are not shallow, thus matching our PTAS for this problem with respect to depth.

1. Introduction

Motivation. Geometric set cover, as the name suggests, is a geometric instantiation of the classical set cover problem. For example, given a set of points and a set of triangles (with fixed locations) in the plane, we want to select a minimum number of triangles that cover all of the given points. Similar geometric variants can be defined for independent set, hitting set, dominating set, and the like.

For relatively simple shapes, such geometric instances should be computationally easier than the general problem. By now there is a large yet incomplete collection of results on such problems, listed below. For example, one can get $(1 + \varepsilon)$ -approximation to the geometric set cover problem when the regions are disks, but we do not have such an approximation algorithm if the regions are fat triangles of similar size. This discrepancy seems arbitrary and somewhat baffling, and the purpose of this work is to better understand these subtle gaps.

*Work on this paper was partially supported by a NSF AF awards CCF-1421231, and CCF-1217462. A preliminary version of this paper appeared in ESA 2015 [HQ15]. A talk by the first author in SODA 2016 was based to some extent on the work in this paper.

[†]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; <http://sarielhp.org/>.

[‡]Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; quanrud2@illinois.edu; <http://illinois.edu/~quanrud2/>.

Objects	Approx. Alg.	Hardness
Disks/pseudo-disks	PTAS [MRR14b], [GRRR16]	Exact version NP-HARD [FG88]
Fat triangles of same size	$O(1)$ [CV07]	APX-HARD: Lemma 55 I.e., no PTAS possible.
Fat objects in \mathbb{R}^2	$O(\log^* \text{opt})$ [ABES14]	APX-HARD: Lemma 55
Objects $\subseteq \mathbb{R}^d$, $O(1)$ density E.g. fat objects, $O(1)$ depth.	PTAS: Theorem 46	Exact version NP-HARD [FG88]
Objects with polylog density	QPTAS: Theorem 46	No PTAS under ETH Lemma 61
Objects with density ρ in \mathbb{R}^d	PTAS: Theorem 46 RT: $n^{O(\rho^{(d+1)/d}/\epsilon^d)}$	No $(1 + \epsilon)$ -approx with RT $n^{\text{poly}(\log \rho, 1/\epsilon)}$ assuming ETH: Lemma 61

Figure 1.1: Known results about the complexity of geometric set-cover. The input consists of a set of points and a set of objects, and the task is to find the smallest subset of objects that covers the points. To see that the hardness proof of Feder and Greene [FG88] indeed implies the above, one just needs to verify that the input instance their proof generates has bounded depth. A QPTAS is an algorithm with running time $n^{O(\text{poly}(\log n, 1/\epsilon))}$.

Plan of attack. We explore the type of graphs that arises out of these geometric instances, and in the process introduce the class of low-density graphs. We explore the properties of this graph class, and the optimization problems that can be approximated efficiently on the broader class of graphs that have small separators. Separability turns out to be the key ingredient needed for efficient approximation. We also study lower bounds for such instances, characterizing when they are computationally hard.

1.1. Background

1.1.1. Optimization problems

Independent set. Given an undirected graph $G = (V, E)$, an *independent set* is a set of vertices $X \subseteq V$ such that no two vertices in X are connected by an edge. It is NP-COMplete to decide if a graph contains an independent set of size k [Kar72], and one cannot approximate the size of the maximum independent set to within a factor of $n^{1-\epsilon}$, for any fixed $\epsilon > 0$, unless $P = NP$ [Hås99].

Dominating set. Given an undirected graph $G = (V, E)$, a *dominating set* is a set of vertices $D \subseteq V$ such that every vertex in G is either in D or adjacent to a vertex in D . It is NP-COMplete to decide if a graph contains a dominating set of size k (by a simple reduction from set cover, which is NP-COMplete [Kar72]), and one cannot obtain a $c \log n$ approximation (for some constant c) unless $P = NP$ [RS97].

1.1.2. Graph classes

Density. Informally, a set of objects in \mathbb{R}^d is *low-density* if no object can intersect too many objects that are larger than it. This notion was introduced by van der Stappen *et al.* [SOBV98], although weaker notions involving a single resolution were studied earlier (e.g. in the work by Schwartz and Sharir [SS85]). *Fatness* is related to density – informally, an object is fat if it contains a ball, and is

contained inside another ball, that up to constant scaling are of the same size. Fat objects have low union complexity [APS08], and in particular, shallow fat objects have low density [Sta92]. Here, a set of shapes is *shallow* if no point is covered by too many of them.

Intersection graphs. A set \mathcal{F} of objects in \mathbb{R}^d induces an *intersection graph* $G_{\mathcal{F}}$ having \mathcal{F} as its set of vertices, and two objects $f, g \in \mathcal{F}$ are connected by an edge if and only if $f \cap g \neq \emptyset$. Without any restrictions, intersection graphs can represent any graph. Motivated by the notion of density, a graph is a *low-density graph* if it can be realized as the intersection graph of a low-density collection of objects in low dimensions.

There is much work on intersection graphs, from interval graphs, to unit disk graphs, and more. The circle packing theorem [Koe36, And70, PA95] implies that every planar graph can be realized as a coin graph, where the vertices are interior disjoint disks, and there is an edge connecting two vertices if their corresponding disks are touching. This implies that planar graphs are low density. Miller *et al.* [MTTV97] studied the intersection graphs of balls (or fat convex object) of bounded depth (i.e., every point is covered by a constant number of balls), and these intersection graphs are readily low density. Some results (weakly) related to our work include: (i) planar graphs are the intersection graph of segments [CG09], and (ii) string graphs (i.e., intersection graph of curves in the plane) have small separators [Mat14].

Polynomial expansion. The class of low-density graphs is contained in the class of graphs with polynomial expansion. This class was defined by Nešetřil and Ossona de Mendez as part of a greater investigation on the sparsity of graphs (see the book [NO12]). A motivating observation to their theory is that the sparsity of a graph (the ratio of edges to vertices) is not necessarily sufficient for tractability. For example, a clique (which is a graph with maximum density) can be disguised as a sparse graph by splitting every edge by a middle vertex. Furthermore, constant degree expanders are also sparse. For both graphs, many optimization problems are intractable (intuitively, because they do not have a small separator).

Graphs with bounded expansion are *nowhere dense graphs* [NO12, Section 5.4]. Grohe *et al.* [GKS14] recently showed that first-order properties are fixed-parameter tractable for nowhere dense graphs. In this paper, we study graphs of polynomial expansion [NO12, Section 5.5], which intuitively requires a graph to not only be sparse, but remain sparse even if one contracts connected patches of vertices of small radii into single vertices. Formally, such locally contracted graphs are called shallow minors – see [Definition 12](#) – and can generally be much denser than the original graph.

It is known that graphs with polynomial expansion have sublinear separators [NO08a]. The converse, that any graph that has hereditary sublinear separators has polynomial expansion, was recently shown by Dvořák and Norin [DN15]. As such, our work looks beyond the geometric setting to consider the general role of separators in approximation.

1.1.3. Further related work

There is a long history of optimization in structured graph classes. Lipton and Tarjan first obtained a PTAS for independent set in planar graphs by using separators [LT79, LT80]. Baker [Bak94] developed techniques for covering problems (e.g. dominated set) on planar graphs. Baker’s approach was extended by Eppstein [Epp00] to graphs with bounded local treewidth, and by Grohe [Gro03] to graphs excluding minors. Separators have also played a key role in geometric optimization algorithms, including: (i) PTAS for independent set and (continuous) piercing set for fat objects [Cha03, MR10], (ii) QPTAS for maximum weighted independent sets of polygons [AW13, AW14, Har14], and (iii) QPTAS for Set Cover by

Objects	Approx. Alg.	Hardness
Disks/pseudo-disks	PTAS [MR10]	Exact version NP-HARD via point-disk duality [FG88]
Fat triangles of similar size	$O(\log \log \text{opt})$ [AES10]	APX-HARD: Lemma 51
Objects with $O(1)$ density	PTAS: Theorem 46	Exact ver. NP-HARD [FG88]
Objects polylog density	QPTAS: Theorem 46	No PTAS under ETH Lemma 61 / Lemma 51
Objects with density ρ in \mathbb{R}^d	PTAS: Theorem 46 Run time $n^{O(\rho^{(d+1)/d}/\epsilon^d)}$	No $(1 + \epsilon)$ -approx with RT $n^{\text{poly}(\log \rho, 1/\epsilon)}$ assuming ETH: Lemma 61

Figure 1.2: Known results about the complexity of *discrete* geometric hitting set. The input is a set of points, and a set of objects, and the task is to find the smallest subset of points such that any object is hit by one of these points.

pseudodisks [MRR14a], among others. Lastly, Cabello and Gajser [CG14a, CG15] develop PTAS’s for some of the problems we study in the specific setting of minor-free graphs.

1.2. Our results

We systematically study the class of graphs that have low density, first proving that they have polynomial expansion. We then develop approximation algorithms for this broader class of graphs, as follows:

- (A) **PTAS for independent set.** For graphs that have sublinear hereditary separators we show PTAS for independent set, see Section 3.1. This covers graphs with low density and polynomial expansion. These results are not surprising in light of known results [CH12], but provide a starting point and contrast for subsequent results.
- (B) **PTAS for packing problems.** The above PTAS also hold for packing problems, such as finding maximal induced planar subgraph, and similar problems, see Example 25 and Lemma 28.
- (C) **PTAS for independent/packing when the output is separable.** More surprisingly, one get a PTAS as long as the subgraph induced on the union of two solutions has polynomial expansion. Thus, while the input may not be sparse, as long as the output is sparse, one can get an efficient approximation algorithms, see Theorem 29.

In particular, this holds if the output is required to have low density, because the union of two sets of objects with low density is still low density. The resulting algorithms in the geometric setting are faster than those for polynomial expansion graphs, by using the underlying geometry of low-density graphs.

- (D) **PTAS for dominating set.** Low density graphs remain low density even if one locally merges objects that are close together, see Lemma 11. More generally, if one considers a collection of t -shallow subgraphs (i.e., subgraphs with radius t in the edge distance) of a polynomial expansion graph, then their intersection graph also has polynomial expansion, as long as no vertex in the original graph participates in more than constant number of subgraphs.

This surprising property implies that local search algorithms provides a PTAS for problems like Dominating Set for graphs with polynomial expansion, see Section 3.3.

- (E) **PTAS for multi-cover dominating set with reach constraints.** These results can be extended to multi-cover variants of dominating set for such graphs, where every vertex can be asked to be dominated a certain number of times, and require that the these dominated vertices are within a certain distance. See [Lemma 41](#).
- (F) **Connected dominating set.** The above algorithms also extend to a PTAS for connected dominating set, see [Section 3.3.6](#).
- (G) **PTAS for vertex cover for graphs with polynomial expansion.** See [Observation 44](#).
- (H) **PTAS for geometric hitting set and set cover.** The new algorithms for dominating sets readily provides PTAS's for discrete geometric set cover and hitting set for low density inputs, see [Section 3.4](#).
- (I) **Hardness of approximation.** The low-density algorithms are complimented by matching hardness results that suggest the new approximation algorithms are nearly optimal with respect to depth (under SETH: the assumption that SAT over n variables can not be solved in better than 2^n time).

The context of our results, for geometric settings, is summarized in [Figure 1.1](#) and [Figure 1.2](#).

Sparsity is not enough. It is natural to hope that the above algorithms would work for sparse graphs (i.e., graphs that have a linear number of edges). Unfortunately, as mentioned earlier, constant degree expanders, which play a prominent role in theoretical computer science, are sparse, and the above algorithms fail for them as they do not have separators.

Low level technical contributions. We show that graphs with polynomial expansion retain polynomial expansion even if one is allowed to locally connect a vertex to other nearby vertices in a controlled way. To this end, we extend the notion of t -shallow minors to shallow packings (see [Definition 33](#)). We then use a probabilistic argument to show that the associated intersection graph still has polynomial expansion, see [Lemma 34](#). The proof of this lemma is elegant and might be of independent interest.

Paper organization. We describe low-density graphs in [Section 2.1](#) and prove some basic properties. Bounded expansion graphs are surveyed in [Section 2.2](#). [Section 3](#) present the new approximation algorithms. [Section 4](#) present the hardness results. Conclusions are provided in [Section 5](#).

2. Preliminaries

2.1. Low-density graphs

Definition 1. For a graph $G = (V, E)$, and any subset $X \subseteq V$, let $G|_X$ denote the *induced subgraph* of G over X . Formally, we have

$$G|_X = (X, \{uv \mid u, v \in X, \text{ and } uv \in E\}).$$

Definition 2. Consider a set of objects \mathcal{U} . The *intersection graph* of \mathcal{U} , denoted by $G_{\mathcal{U}}$, is the graph having \mathcal{U} as its set of vertices, and there is an edge between two objects $f, g \in \mathcal{U}$ if they intersect. Formally,

$$G_{\mathcal{U}} = (\mathcal{U}, \{fg \mid f, g \in \mathcal{U} \text{ and } f \cap g \neq \emptyset\}).$$

One of the two main thrusts of this work is investigating the following family of graphs.

Definition 3. A set of objects \mathcal{U} in \mathbb{R}^d (not necessarily convex or connected) has *density* ρ if any object f (not necessarily in \mathcal{U}) intersects at most ρ objects in \mathcal{U} with diameter greater than or equal to the diameter of f . The minimum such quantity is denoted by $\text{density}(\mathcal{U})$. A graph that can be realized as the intersection graph of a set of objects \mathcal{U} in \mathbb{R}^d with density ρ is *ρ -dense*.

In most of our motivating applications, the density ρ is a constant. We say that a set of objects \mathcal{U} has *low density* if it has constant density $\rho = O(1)$.

When speaking about the density of a graph, we conceptually identify the graph with some set of objects that realize it.

Definition 4. A graph is a *low density graph* if it can be realized as the intersection graph of a set of low density objects in \mathbb{R}^d .

Definition 5. A graph G is *k -degenerate* if any subgraph of G has a vertex of degree at most k .

Lemma 6. A ρ -dense graph $G = (V, E)$ is $(\rho - 1)$ -degenerate, and $|E| \leq (\rho - 1) |V|$.

Proof: Consider the set of objects \mathcal{U} that induces G . Let f be the object with smallest diameter d_0 in \mathcal{U} . By choice of f , any other object intersecting f has diameter at least d_0 . Since at most ρ objects in \mathcal{U} with diameter at least d_0 intersect f (including f itself), the degree of f in G is $\rho - 1$. Clearly, this argument applies to any subgraph of G .

Repeatedly removing the minimum degree vertex from the graph, by the above argument, implies that the number of edges in the graph G is at most $(\rho - 1) |V|$. \blacksquare

2.1.1. Fatness and density

For $\alpha > 0$, an object $g \subseteq \mathbb{R}^d$ is *α -fat* if for any ball b with a center inside g , that does not contain g in its interior, we have $\text{vol}(b \cap g) \geq \alpha \text{vol}(b)$ [BKSV02]¹. A set \mathcal{F} of objects is *fat* if all its members are α -fat for some constant α . A collection of objects \mathcal{U} has *depth* k if any point in the underlying space lies in at most k objects of \mathcal{U} . The depth index of a set of objects is a lower bound on the density of the set, as a point can be viewed as a ball of radius zero. The following is well known, and we include a proof for the sake of completeness – it demonstrates why the notion of density is natural in geometric settings.

Lemma. A set \mathcal{F} of α -fat objects in \mathbb{R}^d with depth k has density $k2^d/\alpha$. In particular, if α, k and d are bounded constants, then \mathcal{F} has low density.

Proof: Let $b = \mathbb{b}(p, r)$ be any ball in \mathbb{R}^d . Let \mathcal{G} be the set of all objects in \mathcal{F} that have diameter larger than b and intersect it, and consider any object $g \in \mathcal{G}$. A ball $\mathbb{b}(q, r)$ centered at a point $q \in g \cap b$ does not contain g in its interior because $\text{diam}(\mathbb{b}(q, r)) \leq \text{diam}(g)$. By the definition of α -fatness, we have

$$\text{vol}(\mathbb{b}(q, r) \cap g) \geq \alpha \text{vol}(\mathbb{b}(q, r)) = \alpha \text{vol}(b).$$

¹There are several different, but roughly equivalent, definitions of fatness in the literature, see de Berg [Ber08] and the followup work by Aronov *et al.* [ABES14] for some recent results. In particular, our definition here is what de Berg refers to as being *locally fat*.

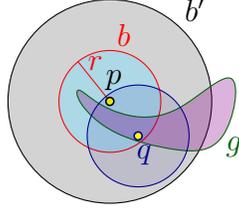


Figure 2.1: Illustration for the proof of [Lemma 1](#).

Furthermore, since $q \in \mathbb{b}(p, r)$, the ball $\mathbb{b}(q, r)$ is contained in the ball $b' = \mathbb{b}(p, 2r)$, and

$$\text{vol}(b' \cap g) \geq \text{vol}(\mathbb{b}(q, r) \cap g) \geq \alpha \text{vol}(b) = \frac{\alpha}{2^d} \text{vol}(b').$$

By assumption, each point in b' can be covered by at most k objects of \mathcal{G} , hence

$$k \text{vol}(b') \geq \sum_{g \in \mathcal{G}} \text{vol}(b' \cap g) \geq |\mathcal{G}| \frac{\alpha}{2^d} \text{vol}(b'). \quad \blacksquare$$

Thus, $|\mathcal{G}| \leq k(2^d/\alpha)$, bounding the number of “large” objects of \mathcal{U} intersecting b .

Definition. A metric space \mathcal{X} is a *doubling space* if there is a universal constant $c_{\text{dbl}} > 0$ such that any ball b of radius r can be covered by c_{dbl} balls of half the radius. Here c_{dbl} is the *doubling constant*, and its logarithm is the *doubling dimension*.

Observation 7. In \mathbb{R}^d the doubling constant is $c_d = 2^{O(d)}$, and the doubling dimension is $O(d)$ [Ver05], making the doubling dimension a natural abstraction of the notion of dimension in the Euclidean case.

The following lemma is used in bounding the density of the intersection graph of shallow minors (i.e., [Lemma 11](#) below).

Lemma 8. Let \mathcal{U} be a set of objects in \mathbb{R}^d with density ρ . Then, for any $\alpha \in (0, 1)$, a ball $b = \mathbb{b}(c, r)$ can intersect at most $\rho c_d^{\lceil \lg 1/\alpha \rceil}$ objects of \mathcal{U} with diameter $\geq 2r\alpha$, where \lg is the log function in base two, and c_d is the doubling constant of \mathbb{R}^d .

Proof: By the definition of the doubling constant, one can cover b by $(c_d)^i$ balls of radius $r/2^i$. As such, one can cover b with $\leq c_d^{\lceil \log_2 1/\alpha \rceil}$ balls of radius $\leq \alpha r$. Each of these balls, by definition of density, can intersect at most ρ objects of \mathcal{U} of diameter at least $2r\alpha$. \blacksquare

2.1.2. Minors of objects

Definition 9. A graph G is *t-shallow* if there is a vertex $h \in V(G)$, such that for any vertex $u \in V(G)$ there is a path π that connects h to u , and π has at most t edges. The vertex h is a *center* of G , denoted by $h = \text{center}(G)$. The minimum integer t such that G is t -shallow is the *radius* of G .

Let \mathcal{U} and \mathcal{V} be two sets of objects in \mathbb{R}^d . The set \mathcal{V} is a *minor* of \mathcal{U} if it can be obtained by deleting objects and replacing pairs of overlapping objects f and g (i.e., $f \cap g \neq \emptyset$) with their union $f \cup g$. Consider a sequence of unions and deletions operations transforming \mathcal{U} into \mathcal{V} . Every object $g \in \mathcal{V}$ corresponds to a set of objects of $C(g) \subseteq \mathcal{U}$, such that $\cup_{h \in C(g)} h = g$. The set $C(g)$ is a *cluster* of objects of \mathcal{U} .

Surprisingly, even for a set \mathcal{F} of fat and convex shapes in the plane with constant density, their intersection graph $G_{\mathcal{F}}$ can have arbitrarily large cliques as minors (see [Figure 2.2](#)). Note that the clusters in [Figure 2.2](#) induce intersection graphs with large graph radius.

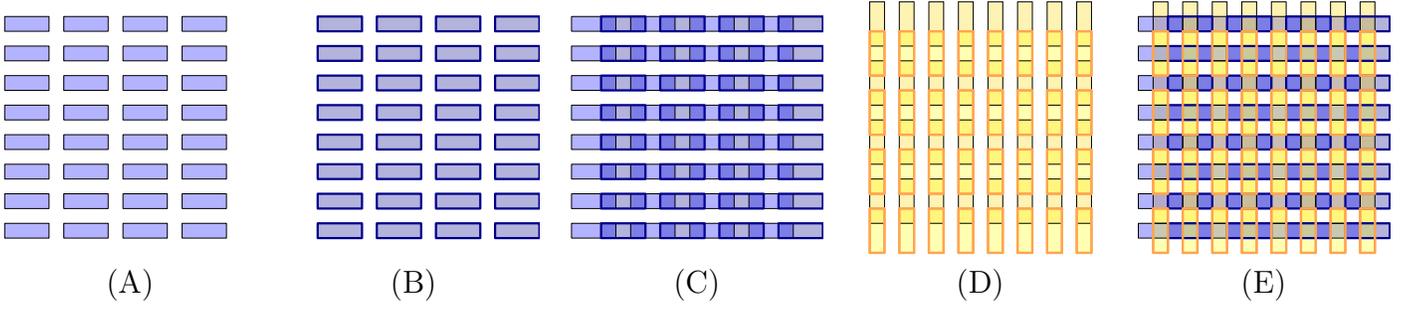


Figure 2.2: (A) and (B) are two low-density collections of n^2 disjoint horizontal slabs, whose intersection graph (C) contains n rows as minors. (D) is the intersection graph of a low-density collection of vertical slabs that contain n columns as minors. In (E), the intersection graph of all the slabs contain the n rows and n columns as minors that form a $K_{n,n}$ bipartite graph, which in turn contains the clique K_n as a minor.

Definition 10. For sets of objects \mathcal{U} and \mathcal{V} , if

- (i) \mathcal{V} is a minor of \mathcal{U} , and
- (ii) the intersection graph of each cluster of \mathcal{U} (that corresponds to an object in \mathcal{V}) is t -shallow,

then \mathcal{V} is a *t -shallow minor* of \mathcal{U} .

The following lemma shows that there is a simple relationship between the depth of a shallow minor of objects and its density.

Lemma 11. *Let \mathcal{U} be a collection of objects with density ρ in \mathbb{R}^d , and let \mathcal{V} be a t -shallow minor of \mathcal{U} . Then \mathcal{V} has density at most $t^{O(d)}\rho$.*

Proof: Every object $g \in \mathcal{V}$ has an associated cluster $C(g) \subseteq \mathcal{U}$ such that $g = \bigcup_{f \in C(g)} f$, and these clusters are disjoint (as subsets of \mathcal{U} ; objects in different may still overlap geometrically). Let $\mathcal{P} = \{C(g) \mid g \in \mathcal{V}\}$ be the induced partition of \mathcal{U} into clusters (which may be a partition of a subset of \mathcal{U}).

Consider any ball $b = \mathfrak{b}(c, r)$, and suppose that $g \in \mathcal{V}$ intersects b and has diameter at least $2r$. Let $C(g) \in \mathcal{P}$ be its cluster, and let $H = G_{C(g)}$ be its associated intersection graph. By assumption, H has (graph) radius at most t , and diameter at most $2t$.

Let h be any object in $C(g)$ that intersect b , let d_H denote the shortest path metric of H (under the number of edges), and let h' be the object in $C(g)$ closest to h (under d_H), such that $\text{diam}(h') \geq r/t$. If there is no such object then $\text{diam}(g) < 2t(r/t) \leq 2r$, which is a contradiction.

Consider the shortest path $\pi \equiv h_1, \dots, h_m$ between $h = h_1$ and $h' = h_m$ in H , where $m \leq 2t$. By the choice of h , we have $\text{diam}(h_i) < r/t$, for $i = 1, \dots, m-1$, and the distance between b and h' is bounded by $\sum_{i=1}^{m-1} \text{diam}(h_i) \leq (m-1)(r/t) < 2r$. The object h' is the *representative* of g , denoted by $\text{rep}(g) \in C(g)$.

Now, let $\mathcal{H} = \{\text{rep}(g) \mid g \in \mathcal{V}, \text{diam}(g) \geq 2r, \text{ and } g \cap b \neq \emptyset\} \subseteq \mathcal{U}$ be the representatives of the large objects in \mathcal{V} intersecting b . The representatives in \mathcal{H} are all distinct, have diameters $\geq r/t$, intersect $\mathfrak{b}(c, 3r)$, and belong to \mathcal{U} - a set with density ρ . Setting $\alpha = 1/6t$, Lemma 8 implies that $|\mathcal{H}| \leq \rho c_d^{\lceil \lg(6t) \rceil}$. Since $c_d = 2^{O(d)}$ [Ver05], it follows that $|\mathcal{H}| = t^{O(d)}\rho$, implying the claim. \blacksquare

2.2. Graphs with polynomial expansion

2.2.1. Basic properties

Definition 12. Let G be an undirected graph. A *minor* of G is a graph H that can be obtained by contracting edges, deleting edges, and deleting vertices from G . If H is a minor of G , then each vertex v of H corresponds to a *cluster* - a connected set $C(v)$ of vertices in G (i.e., these are the vertices of

a tree in the forest formed by the contracted edges). The graph H is a *t -shallow minor* (or a *minor of depth t*) of H , where t is an integer, if for each vertex $v \in V(H)$, the induced subgraph G_C of the corresponding cluster $C = C(v)$ is t -shallow (see [Definition 9](#)). Let $\nabla_t(H)$ denote the set of all graphs that are minors² of H of depth t .

Definition 13 ([NO08a]). The *greatest reduced average density of rank r* , or just the *r -shallow density*, of G is the quantity $\mathfrak{d}_r(G) = \sup_{H \in \nabla_r(G)} \frac{|E(H)|}{|V(H)|}$.

Definition 14. A graph *class* is a (potentially infinite) set of graphs (e.g., the class of planar graphs). The *expansion* of a graph class \mathcal{C} is the function $f : \mathbb{N} \rightarrow \mathbb{N} \cup \{\infty\}$ defined by $f(r) = \sup_{G \in \mathcal{C}} \mathfrak{d}_r(G)$. The class \mathcal{C} has *bounded expansion* if $f(r)$ is finite for all r . Specifically, a class \mathcal{C} with bounded expansion has *polynomial expansion* (resp., *subexponential expansion* or *constant expansion*) if f is bounded by a polynomial (resp., subexponential function or constant). The polynomial expansion is of *order k* if $f(x) = O(x^k)$. Naturally, a graph G has *polynomial expansion* of order k if it belongs to a class of graphs with polynomial expansion of order k .

Observation 15. *If a graph G has bounded expansion, then G has average degree at most $\mu = \mathfrak{d}_1(G)/2 = O(1)$, by taking the graph G as its own 1-shallow minor (with every vertex is its own cluster). In particular, the vertex v_0 with minimum degree has degree at most μ . Similarly, any subgraph of G has a vertex v_1 with degree at most μ , so the graph G , by virtue of its bounded expansion, is $O(1)$ -degenerate (see [Definition 5](#)).*

As an example of a class of graphs with constant expansion, observe that planar graphs have constant expansion because a minor of a planar graph is planar, and by Euler's formula, every planar graph is sparse. More surprisingly, [Lemma 11](#) together with [Lemma 6](#) implies that low-density graphs have polynomial expansion.

Lemma 16. *Let $\rho > 0$ be fixed. The class of ρ -dense graphs in \mathbb{R}^d has polynomial expansion bounded by $f(r) = \rho r^{O(d)}$.*

2.2.2. Separators

Definition 17. Let $G = (V, E)$ be an undirected graph. Two sets $X, Y \subseteq V$ are *separate* in G if (i) X and Y are disjoint, and (ii) there is no edge between the vertices of X and Y in G . A set $Z \subseteq V$ is a *separator* for a set $U \subseteq V$, if $U \setminus Z$ can be partitioned into two *separate* sets X and Y , with $|X| \leq (2/3)|U|$ and $|Y| \leq (2/3)|U|$.

Remark. In the above definition the choice of $2/3$ is arbitrary, and any constant smaller than 1 is sufficient. Furthermore, we are usually interested in separators that are small compared to the domains they separate. Formally, that is $|Z| = o(|U|)$.

Nešetřil and Ossona de Mendez showed that graphs with subexponential expansion have subexponential sized separators. For the simpler case of polynomial expansion, we have the following³.

Theorem 18 ([NO08b, Theorem 8.3]). *Let \mathcal{C} be a class of graphs with polynomial expansion of order k . For any graph $G \in \mathcal{C}$ with n vertices and m edges, one can compute, in $O(mn^{1-\alpha} \log^{1-\alpha} n)$ time, a separator of size $O(n^{1-\alpha} \log^{1-\alpha} n)$, where $\alpha = 1/(2k + 2)$.*

²I.e., these graphs can not legally drink alcohol.

³A proof is also provided in [HQ16].

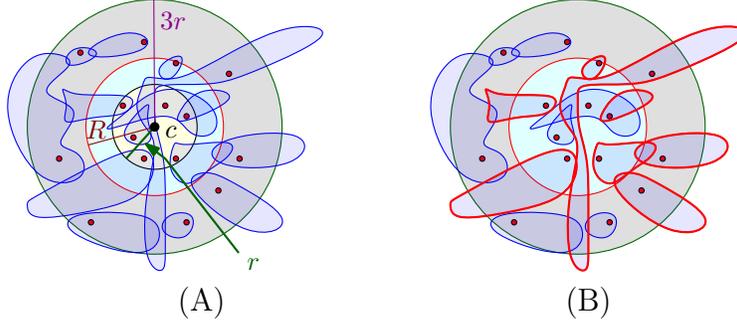


Figure 2.3: Illustration of the proof of [Lemma 19](#). (A) The ball $\mathfrak{b}(c, r)$, and the separating sphere $\mathbb{S}(c, R)$. (B) All the objects intersecting $\mathbb{S}(c, R)$ are in the separating set.

[Theorem 18](#) yields a sublinear separator for ρ -graphs of size

$$O\left((\rho^2 n \log n)^{1 - \frac{1}{O(\log c_{\text{dbl}})}}\right).$$

Geometric arguments give a somewhat stronger separator. For the sake of completeness, we provide next a proof of the following result, but we emphasize that it is essentially already known [[MTTV97](#), [SW98](#), [Cha03](#)]. This proof is arguably simpler and more elegant than previous proofs.

Lemma 19. *Let \mathcal{U} be a set of n objects in \mathbb{R}^d with density $\rho > 0$ (see [Definition 3](#)), and let $k \leq n$ be some prespecified number. Then, one can compute, in expected $O(n)$ time, a sphere \mathbb{S} that intersects $O(\rho^{1/d} k^{1-1/d})$ objects of \mathcal{U} if $k \geq \rho$ (the bound is $O(\rho)$ if $k < \rho$). Furthermore, the number of objects of \mathcal{U} strictly inside \mathbb{S} is at least $k - o(k)$, and at most $O(k)$. For $k = O(n)$ this results in a balanced separator. (Note that the O notation hides constants that depend on d .)*

Proof: For every object $f \in \mathcal{U}$, choose an arbitrary representative point $p_f \in f$. Let P be the resulting set of points. Let $\mathfrak{b}(c, r)$ be the smallest ball containing k points of P . As in [[Har13](#)], randomly pick R uniformly in the range $[r, 2r]$. We claim that the sphere $\mathbb{S} = \mathbb{S}(c, R)$ bounding the ball $b = \mathfrak{b}(c, R)$ is the desired separator.

To this end, consider the distance $\ell = t \cdot r$, where $t \in (0, 1)$ is some real number to be specified shortly. We count the number of objects intersecting \mathbb{S} as follows:

- (A) **Large objects (diameter $\geq \ell$).** The sphere \mathbb{S} can intersect only $N_1 = O(\rho + \rho/t^{d-1})$ objects with diameter $\geq \ell$. Indeed, cover the sphere \mathbb{S} with $O(1/t^{d-1})$ balls of radius $\ell/2$, and let \mathcal{B} be this set of balls. Next, charge each object of diameter larger than ℓ intersecting \mathbb{S} to the ball of \mathcal{B} that intersects it. Each ball of \mathcal{B} get charged at most ρ times.
- (B) **Small objects (diameter $< \ell$).** Let $\mathcal{U}_{b'}$ be the set of objects of \mathcal{U} with diameter $\leq \ell$ fully contained in $b' = \mathfrak{b}(c, 3r)$, which contains all the small objects of \mathcal{U} that intersects \mathbb{S} . The ball b' can be covered by c_d^2 balls of radius r , where c_d is the doubling constant of \mathbb{R}^d (see [Observation 7](#)), and each ball of radius r contains at most k representative points. Thus, b' contains at most $k' = c_d^2 k$ points of P , implying that $|\mathcal{U}_{b'}| \leq k' = O(k)$.

For an object $g \in \mathcal{U}_{b'}$, consider the closest point p and the furthest point q in g from c . The object g is in the separating set (and “intersects” \mathbb{S}) if \mathbb{S} separates p from q (g may be disconnected). As R is chosen uniformly at random from $[r, 2r]$, we have that

$$\alpha(g) = \Pr[g \text{ intersects } \mathbb{S}] \leq \frac{\|c - q\| - \|c - p\|}{r} \leq \frac{\text{diam}(g)}{r} \leq \frac{\ell}{r} = t,$$

by the triangle inequality. The expected number of objects of \mathcal{U}_b intersecting \mathbb{S} is $N_2 = \sum_{g \in \mathcal{U}_b} \alpha(g) \leq tk' = O(tk)$.

We conclude that the separator size, in expectation, is

$$N = N_1 + N_2 = O\left(\rho + \rho/t^{d-1} + kt\right).$$

Solving for $\rho/t^{d-1} = kt$ yields $t = (\rho/k)^{1/d}$, and the resulting separator is (in expectation) of size $N = O(\rho + \rho^{1/d}k^{1-1/d})$.

As for the running time, it is sufficient to find a two approximation to the smallest ball that contains k points of P , and this can be done in linear time [HR15]. Using such an approximation slightly deteriorates the constants in the bounds. By Markov's inequality, \mathbb{S} intersects at most $2N$ objects of \mathcal{U} with probability $\geq 1/2$. If this is not true, we rerun the algorithm. In expectation, the algorithm succeeds in finding a sphere that intersects at most $2N$ objects of \mathcal{U} within a constant number of iterations. ■

Remark. Mark de Berg (personal communication) pointed out the current simplified proof of Lemma 19. The authors thank him for pointing out the simpler proof.

It was recently shown that any graph with strongly sublinear hereditary separators has polynomial expansion [DN15]. In conjunction with the preceding separator (for low-density objects), this yields a second proof that the intersection graphs of low-density objects have polynomial expansion, however with weaker bounds.

A weighted version of the above separator follows by a similar argument.

Lemma 20. *Let \mathcal{U} be a set of n objects in \mathbb{R}^d with density ρ , and weights $w : \mathcal{U} \rightarrow \mathbb{R}$. Let $W = \sum_{f \in \mathcal{U}} w(f)$ be the total weight of all objects in \mathcal{U} . Then one can compute, in expected linear time, a sphere \mathbb{S} that intersects $O(\rho + \rho^{1/d}n^{1-1/d})$ objects of \mathcal{U} . Furthermore, the total weight of objects of \mathcal{U} strictly inside/outside \mathbb{S} is at most cW , where c is a constant that depends only on d .*

Proof: The argument follows the one used in Lemma 19. We pick a representative point from each object, and assign it the weight of the object. Next, we compute a two approximation to the smallest ball containing $\geq cW$ of the total weight of the points (using the algorithm of [HR15] which works also for the weighted version), and the rest of the proof follows readily, observing that in the worst case, n objects might be involved in the calculations. ■

2.2.3. Divisions

Consider a set V . A **cover** of V is a set $\mathcal{W} = \{C_i \subseteq V \mid i = 1, \dots, k\}$ such that $V = \bigcup_{i=1}^k C_i$. A set $C_i \in \mathcal{W}$ is a **cluster**. A cover of a graph $G = (V, E)$ is a cover of its vertices. Given a cover \mathcal{W} , the **excess** of a vertex $v \in V$ that appears in j clusters is $j - 1$. The **total excess** of the cover \mathcal{W} is the sum of excesses over all vertices in V .

Definition 21. A cover \mathcal{C} of G is a **λ -division** if (i) for any two clusters $C, C' \in \mathcal{C}$, the sets $C \setminus C'$ and $C' \setminus C$ are separated in G (i.e., there is no edge between these sets of vertices in G), and (ii) for all clusters $C \in \mathcal{C}$, we have $|C| \leq \lambda$.

A vertex $v \in V$ is an **interior vertex** of a cover \mathcal{W} if it appears in exactly one cluster of \mathcal{W} (and its excess is zero), and a **boundary vertex** otherwise. By property (i), the entire neighborhood of an interior vertex of a division lies in the same cluster.

Remark. A division is not only a cover of the vertices, but also a cover of the edges. Consider a λ -division \mathcal{C} of a graph G , and an edge $uv \in E(G)$. We claim that there must be a cluster C in \mathcal{C} , such that both u and v are in C . Indeed, if not, then there are two clusters C_u and C_v , such that $u \in C_u$ and $v \in C_v$, but then $C_u \setminus C_v \ni u$ and $C_v \setminus C_u \ni v$ are not separated in G , contradicting the definition.

Remark. The property of having λ -divisions is slightly stronger than being weakly hyperfinite. Specifically, a graph is *weakly hyperfinite* if there is a small subset of vertices whose removal leaves small connected components [NO12, Section 16.2]. Clearly, λ -divisions also provide such a set (i.e., the boundary vertices). The connected components induced by removing the boundary vertices are not only small, but the neighborhoods of these components are small as well.

As noted by Henzinger *et al.* [HKRS97], strongly sublinear separators obtain λ -divisions with total excess εn for $\lambda = \text{poly}(1/\varepsilon)$. Such divisions were first used by Frederickson in planar graphs [Fre87]. In general, the stronger the separator (i.e., the smaller the separating sets), the less excess is generated in obtaining λ -divisions for some prescribed λ . The following known result quantifies the relationship between the size of the separator and the excess generated, and a full proof is provided in [HQ16].

Lemma 22. *Let G be a graph with n vertices, such that any induced subgraph with m vertices has a separator with $O(m^\alpha \log^\beta m)$ vertices, for some $\alpha < 1$ and $\beta \geq 0$. Then, for $\varepsilon > 0$, the graph G has λ -divisions with total excess εn , where $\lambda = O\left((\varepsilon^{-1} \log^\beta \varepsilon^{-1})^{1/(1-\alpha)}\right)$. Furthermore, if a balanced separator can be found in polynomial time, then the λ -division can be computed in polynomial time.*

Remark 23 (Divisions for weaker separators). One can still obtain divisions for graph classes with weaker separators of size $O\left(n/\log^{O(1)} n\right)$. Rather than a $\text{poly}(1/\varepsilon)$ -division with excess εn , we get a $f(1/\varepsilon)$ -division with excess εn for some function f . Consequently, the PTAS throughout this paper extend to a slightly broader class of graphs than polynomial expansion, see [Remark 45](#).

Corollary 24. (A) *Let G be a graph with polynomial expansion of degree k and n vertices, and let $\varepsilon > 0$ be fixed. Then G has $O\left((1/\varepsilon)^{2k+2} \log^{2k+1}(1/\varepsilon)\right)$ -divisions with total excess εn .*

(B) *Let $G = (V, E)$ be a ρ -dense graph with n vertices arising out of a given set of objects in \mathbb{R}^d . Then G has λ -divisions, with $\lambda = O(\rho/\varepsilon^d)$ and total excess at most εn . This division can be computed in expected $O(n \log(n/\lambda))$ time.*

Proof: (A) By [Theorem 18](#), G has separator with parameters $\alpha = 1 - 1/(2k + 2)$ and $\beta = 1 - 1/(2k + 2)$. Plugging this into [Lemma 22](#) implies λ -divisions where $\lambda = O\left(\left((1/\varepsilon) \log^\beta(1/\varepsilon)\right)^{1/(1-\alpha)}\right) = O\left((1/\varepsilon)^{2k+2} \log^{2k+1}(1/\varepsilon)\right)$.

(B) By [Lemma 19](#), any subgraph of G with m vertices has a separator of size $\leq c(\rho + \rho^{1/d} m^{1-1/d})$, for some constant c . One can break up G in the natural recursive fashion using separators (see the proof of [Lemma 22](#) in [HQ16] for details), until each portion has size m , and $c(\rho + \rho^{1/d} m^{1-1/d}) \leq \varepsilon m/c'$, where c' is some absolute constant. As can be easily verified, this holds for $m = \Omega(\rho/\varepsilon^d)$. Setting $\lambda = m$ implies that the resulting λ -divisions with excess $\leq \varepsilon n$.

As for the running time, computing the separator for a graph with m vertices takes expected $O(m)$ time (assuming basic operation like deciding if an object intersects a sphere can be done in constant time), using the algorithm of [Lemma 19](#), and the recursion depth is $O(\log(n/\lambda))$. ■

2.3. Hereditary and mergeable properties

Let $\Pi \subseteq 2^V$ be a property defined over subsets of vertices of a graph $G = (V, E)$ (e.g., Π is the set of all independent sets of vertices in G). The property Π is *hereditary* if for any $X \subseteq Y \subseteq V$, if Y satisfies Π , then X satisfies Π . The property Π is *mergeable* if for any $X, Y \subseteq V$ that are separate in G , if X and Y each satisfy Π , then $X \cup Y$ satisfies Π . We assume that whether or not $X \in \Pi$ can be checked in polynomial time.

Given a set \mathcal{F} and a property $\Pi \subseteq 2^{\mathcal{F}}$, the *packing problem* associated with Π , asks to find the largest subset of \mathcal{F} satisfying Π .

Example 25. Some geometric flavors of packing problems that corresponds to hereditary and mergeable properties include:

- (A) Given a collection of objects \mathcal{U} , find a maximum independent subset of \mathcal{U} .
- (B) Given a collection of objects \mathcal{U} , find a maximum subset of \mathcal{U} with density at most ρ , where ρ is prespecified.
- (C) Find a maximum subset of \mathcal{U} whose intersection graph is planar or otherwise excludes a prespecified graph minor.
- (D) Given a point set P , a constant k , and a collection of objects \mathcal{U} , find the maximum subset of \mathcal{U} such that each point in P is contained in at most k objects in \mathcal{U} .

3. Approximation algorithms

3.1. Approximation algorithms using separators

Graphs whose induced subgraphs have sublinear and efficiently computable separators are already strong enough to yield PTAS for mergeable and hereditary properties (see [Section 2.3](#) for relevant definitions). Such algorithms are relatively easy to derive, and we describe them as a contrast to subsequent results, where such an approach no longer works. As the following testifies, one can $(1 - \varepsilon)$ -approximate, in polynomial time, the independent set in a low-density or polynomial-expansion graphs (as independent set is a mergeable and hereditary property). The basic idea appears in [\[AST90\]](#), and we include a proof for the sake of completeness.

Lemma 26. *Let $G = (V, E)$ be a graph with n vertices, with the following properties:*

- (A) *Any induced subgraph of G on m vertices has a separator with $O(m^\alpha \log^\beta m)$ vertices, for some constants $\alpha < 1$ and $\beta \geq 0$, and this separator can be computed in polynomial time. (I.e., low density and polynomial expansion graphs have such separators.)*
- (B) *There is a hereditary and mergeable property Π defined over subsets of vertices of G .*
- (C) *The largest set $O \in \Pi$, is of size at least n/c , where c is some absolute constant.*

Then, for any $\varepsilon > 0$, one can compute, in $O(n^{O(1)} + 2^\lambda \lambda^{O(1)} n)$ time, a set $X \in \Pi$ such that $|X| \geq (1 - \varepsilon)|O|$, where $\lambda = O\left((\varepsilon^{-1} \log^\beta \varepsilon^{-1})^{1/(1-\alpha)}\right)$.

Proof: Set $\delta = \varepsilon/2c$. By [Lemma 22](#), one can compute a λ -division for G in polynomial time, such that its total excess is $\mathcal{E} \leq \delta n \leq \varepsilon n/2c \leq \varepsilon |O|/2$, where λ is as stated above. Throw away all the boundary vertices of this division, which discards at most $2\mathcal{E} \leq \varepsilon |O|$ vertices. The remaining clusters are separated from one another, and have size at most λ . For each cluster, we can find its largest subset with property Π by brute force enumeration in $O(2^\lambda \lambda^{O(1)})$ time per cluster. Then we merge the sets computed for each cluster to get the overall solution. Clearly, the size of the merged set is at least $|O| - 2\mathcal{E} \geq (1 - \varepsilon)|O|$. The overall running time of the algorithm is $O(n^{O(1)} + 2^\lambda \lambda^{O(1)} n)$. ■

Example 27 (Largest induced planar subgraph). Consider a graph $G = (V, E)$ with n vertices and with bounded expansion of order k . Assume, that the task is to find the largest subset $X \subseteq V$, such that the induced subgraph $G|_X$ is, say, a planar graph. Clearly, this property is hereditary and mergeable, and checking if a specific induced subgraph is planar can be done in linear time [HT74].

By **Observation 15**, the graph G is t -degenerate, for some $t = O(1)$, since G has bounded expansion. Consequently, G contains an independent set of size $\geq n/(t+1) = \Omega(n)$. This independent set is a valid induced planar subgraph of size $\Omega(n)$. Thus, the algorithm of **Lemma 26** applies, resulting in an $(1 - \varepsilon)$ -approximation to the largest induced planar subgraph. The running time of the resulting algorithm is $n^{O(1)} + f(k, \varepsilon)n$, for some function f .

Lemma 28. *Let $\varepsilon > 0$ be a parameter, and \mathcal{U} be a given set of n objects in \mathbb{R}^d that are ρ -dense. Then one can compute a $(1 - \varepsilon)$ -approximation to the largest independent set in \mathcal{U} . The running time of the algorithm is $O(n \log n + 2^\lambda \lambda^{O(1)} n)$, where $\lambda = O(\rho^{d+1}/\varepsilon^d)$.*

*More generally, one can compute, with the same running time, a $(1 - \varepsilon)$ -approximate solution for all the problems described in **Example 25**.*

Proof: Consider the intersection graph $G = G_{\mathcal{U}}$, and observe that by the low-density property, it always have a vertex of degree $< \rho$ (i.e., take the object in \mathcal{U} with the smallest diameter). As such, removing this object and its neighbors from the graph, adding it to the independent set and repeating this process, results in an independent set in G of size n/ρ . Thus implying that the largest independent set has size $\Omega(n)$. Now, apply the algorithm of **Lemma 26** to G using the improved λ -divisions of **Corollary 24 (B)**. Here, we need the total excess to be bounded by $(\varepsilon/\rho)n$, which implies that $\lambda = O(\rho/(\varepsilon/\rho)^d) = O(\rho^{d+1}/\varepsilon^d)$.

For the second part, observe that all the problems mentioned in **Example 25** have solutions bigger than the independent set of \mathcal{U} , and the same algorithm applies with minor modifications. ■

Remark. For computing the largest independent set, one does not need to assume the low density on the input – a more elaborate algorithm works, see **Lemma 31** below.

It is tempting to try and solve problems like dominating set on polynomial-expansion graphs using the algorithm of **Lemma 28**. However, note that a dominating set in such a graph (or even in a star graph) might be arbitrarily smaller than the size of the graph. Thus, having small divisions is not enough for such problems, and one needs some additional structure.

3.2. Local search for independent set and packing problems

Chan and Har-Peled [CH12] gave a PTAS for independent set with planar graphs, and the algorithm and its underlying argument extends to hereditary graph classes with strongly sublinear separators (see also the work by Mustafa and Ray [MR10]).

3.2.1. Definitions

Let Π be a hereditary and mergeable property, and let λ be a fixed integer. For two sets, X and Y , their *symmetric difference* is $X \Delta Y = (X \setminus Y) \cup (Y \setminus X)$. Two vertex sets X and Y are **λ -close** if $|X \Delta Y| \leq \lambda$; that is, if one can transform X into Y by adding and removing at most λ vertices from X . A vertex set $X \in \Pi$ is **λ -locally optimal** in Π if there is no $Y \in \Pi$ that is λ -close to X and “improves” upon X . In a maximization problem Y **improves** $X \iff |Y| > |X|$. In a minimization problem, an improvement decreases the cardinality.

3.2.2. The local search algorithm in detail

The λ -local search algorithm starts with an arbitrary (and potentially empty) solution $X \in \Pi$ and, by examining all λ -close sets, repeatedly makes λ -close improvements until terminating at a λ -locally optimal solution. Each improvement in a maximization (resp., minimization) problem increases (resp., decreases) the cardinality of the set, so there are at most n rounds of improvements, where n is the size of the ground set of Π . Within a round we can exhaustively try all exchanges in time $n^{O(\lambda)}$, bounding the total running time by $n^{O(\lambda)}$.

3.2.3. Analysis of the algorithm

Theorem 29. *Let $G = (V, E)$ be a given graph with n vertices, and let Π be a hereditary and mergeable property defined over the vertices of G that can be tested in polynomial time. Furthermore, let $\varepsilon > 0$ and λ be parameters, and assume that for any two sets $X, Y \subseteq V$, such that $X, Y \in \Pi$, we have that $G_{|X \cup Y}$ has a λ -division with total excess $\varepsilon |X \cup Y|$. Then, the λ -local search algorithm computes, in $n^{O(\lambda)}$ time, a $(1 - 2\varepsilon)$ -approximation to the maximum size set $Z \subseteq V$ satisfying $Z \in \Pi$.*

Proof: Let $O \subseteq V$ be an optimal maximum set satisfying Π , and L be a λ -locally maximal set satisfying Π . Consider the induced subgraph $K = G_{O \cup L}$, and observe that, by assumption, there exists a λ -division $\mathcal{W} = \{C_1, \dots, C_m\}$ of K , with boundary vertices B and $\text{excess}(\mathcal{W}) \leq \varepsilon |O \cup L| \leq 2\varepsilon |O|$. For $i = 1, \dots, m$, let

$$\begin{aligned} O_i &= (O \cap C_i) \setminus B, & o_i &= |O_i|, \\ L_i &= (L \cap C_i), & l_i &= |L_i|, \\ B_i &= B \cap C_i, & \text{and } b_i &= |B_i|. \end{aligned}$$

Fix i , and consider the set $L' = (L \setminus L_i) \cup O_i$. Since $|C_i| \leq \lambda$, the set L' is λ -close to L . Since Π is hereditary, $L \setminus L_i \in \Pi$, and since O_i and $L \setminus L_i$ are separated, their union L' is in Π . Thus, the local search algorithm considers the valid exchange from L to L' . As the set L is λ -locally optimal, the exchange replacing L_i by O_i can not increase the overall cardinality. Since $|L| - l_i + o_i = |L'| \leq |L|$, this implies that $l_i \geq o_i$. Summing over all i , we have

$$|L| \geq \sum_{i=1}^m l_i - \sum_{i=1}^m b_i \geq \sum_{i=1}^m o_i - \sum_{i=1}^m b_i \geq |O| - \text{excess}(\mathcal{W}) \geq (1 - 2\varepsilon) |O|,$$

as desired. ■

Remark 30. It is illuminating to consider the requirements to make the argument of [Theorem 29](#) go through. We need to be able to break up the conflict graph between the local and optimal solutions into small clusters, such that the total number of boundary vertices (counted with repetition) is small. Surprisingly, even if all (or most of) the vertices of a single cluster are boundary vertices, the argument still goes through.

Lemma 31. *Let $\varepsilon > 0$ and ρ be parameters, and let \mathcal{U} be a given collection of objects in \mathbb{R}^d such that any independent set in \mathcal{U} has density ρ . Then the local search algorithm computes a $(1 - \varepsilon)$ -approximation to the maximum size independent subset of \mathcal{U} in time $n^{O(\rho/\varepsilon^{d+1})}$.*

Proof: Observing that the union of two ρ -dense sets results in a 2ρ -dense set, and using the algorithm of [Theorem 29](#), together with the divisions of [Corollary 24 \(B\)](#), implies the result. ■

Remark. (A) We emphasize that [Lemma 31](#) requires only that independent sets of the input objects \mathcal{U} have low density – the overall set \mathcal{U} might have arbitrarily large density.

(B) All the problems of [Example 25](#) have a PTAS using the [Lemma 31](#) as long as the output has low density.

3.3. Dominating Set

We are interested in approximation algorithms for the following generalization of the dominating set problem.

Definition 32. In the *dominating subset problem*, one is given an undirected graph $G = (V, E)$ and two subsets of vertices R and D , such that D dominates R . The set D *dominates* R if every vertex in R either is in D or is adjacent to some vertex in D . The task is to compute the smallest subset of D that still dominates R .

One can approximate the dominated set, as the reader might expect, via a local search algorithm. Before analyzing the algorithm, we need to develop some tools to be able to argue about the interaction between the local and optimal solution.

3.3.1. Shallow packings

Definition 33. Given a graph $G = (V, E)$, a collection of sets $\mathcal{F} = \{C_i \subseteq V \mid i = 1, \dots, t\}$ is a (ω, t) -*shallow packing* of G , or just a (ω, t) -*packing*, if for all i , the induced graph $G|_{C_i}$ is t -shallow (see [Definition 9](#)), and every vertex of V appears in at most ω sets of \mathcal{F} ⁴.

The *induced packing graph* $G[\mathcal{F}]$ has \mathcal{F} as the set of vertices, and two clusters $C, C' \in \mathcal{F}$ are connected by an edge if they share a vertex (i.e., $C \cap C' \neq \emptyset$), or there are vertices $u \in C$ and $v \in C'$, such that $uv \in E$.

For example, the induced packing graph of a $(1, t)$ -packing is the t -shallow minor induced by the clusters of the packing (see [Definition 12](#)).

In the following lemma, we use a scheme that is similar to algorithms for computing random partitions of metric spaces [[CKR01](#), [FRT04](#), [MS09](#)], nets in sparse graphs [[EHS15](#)], and is also similar in spirit to randomized incremental construction of Voronoi diagrams [[KMM93](#)].

Lemma 34. *Let $G = (V, E)$ be an undirected graph, and \mathcal{F} an (ω, t) -packing of G . Then the induced packing graph $H = G[\mathcal{F}]$ has edge density $\frac{|E(H)|}{|V(H)|} \leq 2(t+1)^2\omega^2 d_t(G) + \omega$, where $d_t(G)$ is the t -shallow density of G , see [Definition 13](#).*

Proof: Let the clusters of \mathcal{F} be $\{C_1, \dots, C_m\}$. For each cluster $C_i \in \mathcal{F}$, designate a center vertex $c_i \in C_i$ that can reach any other vertex in C_i by a path contained in C_i of length t or less. Let $\pi : [m] \rightarrow [m]$ be a random permutation of the cluster indices, chosen uniformly at random, and initialize $\mathcal{F}' = \emptyset$, where $[m] = \{1, \dots, m\}$. For $i = 1, \dots, m$, in order, check if $c_{\pi(i)}$ has been “scooped”; that is, if $c_{\pi(i)} \in \bigcup_{C' \in \mathcal{F}'} C'$, and if so, ignore it. Otherwise, let $C'_{\pi(i)}$ be the set of vertices of the connected component of $c_{\pi(i)}$ in the induced subgraph of G over $C_{\pi(i)} \setminus \bigcup_{C' \in \mathcal{F}'} C'$, and add $C'_{\pi(i)}$ to \mathcal{F}' . Intuitively, the set \mathcal{F}' is a $(1, t)$ -packing of G resulting from randomly shrinking the clusters of \mathcal{F} .

⁴We allow a set C to appear in \mathcal{F} more than once; that is, \mathcal{F} is a multiset. Also, this definition is unrelated to the notion of shallow packings in [[DEG16](#)].

We bound the number of edges in $H = G[\mathcal{F}]$ by a function of the expected number of edges in the random graph $H' = G[\mathcal{F}']$. Let $E_1 = \{C_i C_j \in E(H) \mid c_i \in C_j \text{ or } c_j \in C_i\}$ be the set of edges between pairs of clusters where the center of one cluster is also in the other cluster. Since a center c_i can be covered at most ω times by \mathcal{F} , we have $|E_1| \leq \omega |\mathcal{F}|$. Next, consider the set of remaining edges,

$$E_2 = \{C_i C_j \in E(H) \mid c_j \notin C_i \text{ and } c_i \notin C_j\},$$

between adjacent clusters where neither center lies in the opposing cluster. For an edge $C_i C_j \in E_2$, consider the probability that $C'_i C'_j \in E(H')$.

Since C_i and C_j are adjacent in H , there is a path P in G from c_i to c_j of length at most $2t + 1$ that is contained in $C_i \cup C_j$, and a sufficient condition for $C'_i C'_j \in E(H')$ is that P is contained in $C'_i \cup C'_j$. This holds if the permutation π ranks i and j ahead of any other index k such that C_k intersects the vertices of P . There are at most $2t + 2$ vertices on P , where each vertex can appear in at most ω clusters of \mathcal{F} , and overall there are at most

$$\ell = 2(t + 1)\omega$$

clusters that compete for control over the vertices of P in \mathcal{F}' . The probability that, among these relevant clusters, the random permutation π ranks i and j before all others is $2(\ell - 2)!/\ell! \geq 2/\ell^2$. Therefore, for $C_i C_j \in E_2$, we have $\Pr[C'_i C'_j \in E(H')] \geq 2/\ell^2$. By linearity of expectation, and since $H' = G[\mathcal{F}']$ is a t -shallow minor of G , we have

$$\begin{aligned} |E_2| &= \sum_{e \in E_2} \frac{\ell^2/2}{\ell^2/2} \leq \frac{\ell^2}{2} \sum_{e \in E_2} \Pr[e \in E(H')] = \frac{\ell^2}{2} \mathbf{E}[|E(H')|] \leq \frac{\ell^2}{2} \mathfrak{d}_t(G) |\mathcal{F}'| \\ &\leq \frac{\ell^2}{2} \mathfrak{d}_t(G) |\mathcal{F}|. \end{aligned}$$

We conclude that $\frac{|E(H)|}{|V(H)|} = \frac{|E_2|}{|\mathcal{F}|} + \frac{|E_1|}{|\mathcal{F}|} \leq (\ell^2/2)\mathfrak{d}_t(G) + \omega$, as desired. \blacksquare

Remark. Results similar to [Lemma 34](#) are already known [[NO12](#)]. However, our result has a polynomial dependency on ω and t , while the known results seems to imply an exponential dependency.

[Lemma 34](#) bounds the density of $G[\mathcal{F}]$, and we next bound the u -shallow density of $G[\mathcal{F}]$.

Lemma 35. *Consider a graph G , and an (ω, t) -packing \mathcal{F} of G . Then, for any integer $u > 0$, we have*

$$\mathfrak{d}_u(G[\mathcal{F}]) \leq \omega^2(2u + 1)^2(2t + 1)^2 \mathfrak{d}_{2tu+t+u}(G) + \omega.$$

In particular, if t and ω are constants, and G has polynomial expansion of order k , then $G[\mathcal{F}]$ has polynomial expansion of order $k + 2$.

Proof: For $u \geq 1$, consider a u -shallow minor H of $G[\mathcal{F}]$. Every cluster in this cover corresponds to an expanded cluster in the original graph G with radius at most $2tu + u + t$, bounded as follows. Fix a cluster C with center c (as a subgraph of $G[\mathcal{F}]$, that contracts to a vertex in H). Each vertex $v \in C$ to a cluster D_v in the first graph G with radius at most t . As a subgraph of G , C expands to the induced subgraph over $D = \bigcup_{v \in C} D_v$.

Let $d \in D$ be the center of the ‘‘center cluster’’ D_c . We claim that d can reach any other vertex z in D by a path of at most $2tu + u + t$ edges. Indeed, let $z \in D_v$ for some cluster $v \in C$. Since C has radius u , there exists a path $v_0 = c, v_1, \dots, v_{u-1}, v_u = v$ of $u + 1$ vertices from c to v in $G[\mathcal{F}]$. In G ,

this path expands to $u + 1$ clusters $D_{v_0} = D_c, D_{v_1}, D_{v_2}, D_{v_3}, \dots, D_{v_{u-1}}, D_u$, where each two consecutive clusters are connected by an edge. For $i = 1, \dots, u$, let $s_i t_i$ be an edge from $D_{v_{i-1}}$ to D_{v_i} . Let p_0 be a path from d to s_1 inside D_d of length $\leq t$, p_1 be a path from t_u to z inside D_{v_u} length at most $2t$, and for $i = 1, \dots, u - 1$, let p_i be a path of length at most $2t$ from t_i to s_{i+1} inside D_{v_i} . Then the concatenated path

$$c \xrightarrow{p_0} s_1 \rightarrow t_1 \xrightarrow{p_1} s_2 \rightarrow \dots \rightarrow t_{u-1} \xrightarrow{p_{u-1}} s_u \rightarrow t_u \xrightarrow{p_u} z$$

is a path of c to z of length at most $u + t + 2tu$, as desired.

Now, a vertex in G might participate in ω clusters of the shallow minor H . That is, the resulting set \mathcal{G} of clusters is an $(\omega, 2tu + u + t)$ -packing of G . By [Lemma 34](#), we have

$$\begin{aligned} \alpha(H) &= \frac{|E(H)|}{|V(H)|} = \frac{|E(G[\mathcal{G}])|}{|V(G[\mathcal{G}])|} \leq 2\omega^2(2tu + u + t + 1)^2 \mathfrak{d}_{2tu+t+u}(G) + \omega \\ &= 2\omega^2(t + 1)^2(u + 1)^2 \mathfrak{d}_{2tu+t+u}(G) + \omega. \end{aligned}$$

By [Definition 13](#), we have

$$\mathfrak{d}_u(G[\mathcal{F}]) = \sup_{H \in \nabla_u(G[\mathcal{F}])} \alpha(H) \leq 2\omega^2(t + 1)^2(u + 1)^2 \mathfrak{d}_{2tu+t+u}(G) + \omega,$$

as desired. ■

3.3.2. Lexicographic product and shallow density

An interesting consequence of the above is an improvement over known bounds for the shallow density under lexicographic product (this result is not required for the rest of the paper). Given two graphs G and H , the *lexicographic product* $G \bullet H$ is the graph obtained by blowing up each vertex in G with a copy of H . More formally, $G \bullet H$ has vertex set $V(G) \times V(H)$ and an edge between two vertices (x, y) and (x', y') if either (a) $xx' \in E(G)$, or (b) $x = x'$ and $yy' \in E(H)$.

Corollary 36. *For any graph G , clique K_ω , and $t \in \mathbb{N}$, we have $\mathfrak{d}_t(G \bullet K_\omega) \leq 2\omega^2(t + 1)^2 \mathfrak{d}_t(G) + \omega$. In particular, if ω is constant and G has polynomial expansion of order k , then $G \bullet K_\omega$ has polynomial expansion of order $k + 2$.*

Proof: A t -shallow minor of $G \bullet K_\omega$ is the induced packing graph of the (ω, t) -packing in G formed by projecting each cluster onto its first coordinates in G . Thus, the claimed inequality follows from [Lemma 34](#). ■

[Corollary 36](#) is an exponential improvement over the best previously known bounds, on the order of $\mathfrak{d}_t(G \bullet K_\omega) \leq \left(O(\omega t \mathfrak{d}_t(G))\right)^{O(t)}$, by Nešetřil and Ossona de Mendez [[NO08a](#)] (see also the comments following the proof of [Proposition 4.6](#) in [[NO12](#)]).

3.3.3. Low density objects and (ω, t) -packings

Definition 37. For a set of objects \mathcal{U} , a collection of subsets $\mathcal{F} = \{C_i \subseteq \mathcal{U} \mid i = 1, \dots, t\}$ forms a *(ω, t) -shallow packing* of G if, for all i , the intersection graph G_{C_i} is t -shallow (see [Definition 9](#)), and every object of \mathcal{U} appears in at most ω sets of \mathcal{F} . The *induced object set* $\mathcal{U}[\mathcal{F}]$ is the collection of objects $\left\{ \bigcup_{f \in C_i} f \mid C_i \in \mathcal{F} \right\}$ formed by taking the union of each cluster in \mathcal{F} .

Lemma 38. *Let \mathcal{U} be a collection of objects with density ρ in \mathbb{R}^d , and let \mathcal{F} be an (ω, t) -shallow packing. Then the induced object set $\mathcal{U}[\mathcal{F}]$ has density $O(\omega\rho t^{O(d)})$.*

Proof: Consider the multiset of objects $\mathcal{V} = \bigcup_{C_i \in \mathcal{F}} C_i$ where each object $f \in \mathcal{U}$ is repeated according to its multiplicity in \mathcal{F} . Since each object in \mathcal{U} appears in \mathcal{V} at most ω times, \mathcal{V} has density $\omega\rho$. Every cluster $C \in \mathcal{F}$ can be interpreted as a new cluster $f(C)$ of objects of \mathcal{V} , where the resulting set of clusters $\mathcal{F}' = \{f(C) \mid C \in \mathcal{F}\}$ are now disjoint.

As such, $\mathcal{V}[\mathcal{F}']$ is a t -shallow minor of \mathcal{V} . By [Lemma 11](#), the graph $\mathcal{V}[\mathcal{F}']$ has density $O(\omega\rho t^{O(d)})$. ■

3.3.4. The result

Shallow packings arise in the analysis of the approximation algorithm for dominating set, where vertices are clustered together around the vertex that dominates them. In this setting, we prefer the following simple and convenient terminology.

Definition 39. Given a dominating set $D = \{v_1, \dots, v_m\}$ of vertices in a graph $G = (V, E)$, and a set of vertices $R \subseteq V$ being dominated by D , we generate a sequence of clusters $C_1, \dots, C_m \subseteq D \cup R$ that specifies for every element of D , which elements it covers.

Initially, we set $D_0 = D$ and $R_0 = R$. In the i th iteration, for $i = 1, \dots, m$, let

$$C_i = \{v_i\} \cup ((N(v_i) \cap R_{i-1}) \setminus D_{i-1}), \quad D_i = D_{i-1} \setminus \{v_i\}, \quad \text{and } R_i = R_{i-1} \setminus C_i,$$

where $N(v_i)$ is the set of vertices adjacent to v_i in G . Conceptually, C_i induces a star-like graph G_i over C_i , where every vertex of C_i is connected to v_i . The cluster C_i (and implicitly to G_i) is a **flower**, where v_i is its **head**. The collection of clusters $\mathfrak{F}(D, R) = \{C_1, \dots, C_m\}$ is the **flower decomposition** of the given instance. Note that a flower is a 1-shallow graph, and a flower decomposition is a $(1, 1)$ -shallow packing.

Theorem 40. *Let $G = (V, E)$ be a graph with n vertices and with polynomial expansion of order k , let $R, D \subseteq V$ be two sets of vertices such that D dominates R , and let $\varepsilon > 0$ be fixed. Then, for $\lambda = O(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon))$, the λ -local search algorithm computes, in $n^{O(\lambda)}$ time, a $(1 + \varepsilon)$ -approximation to the smallest cardinality subset of D that dominates R .*

Proof: The algorithm starts with the whole collection D as the local solution, and performs legal local exchanges of size λ that decrease the size of the local solution by at least one until no such exchange is available (see [Section 3.2.2](#)).

Let $O \subseteq D$ and $L \subseteq D$ be the optimal and locally minimal sets dominating R , respectively. Let $\mathcal{O} = \mathfrak{F}(O, R)$ and $\mathcal{L} = \mathfrak{F}(L, R)$ be the corresponding flower decompositions. In the following, for vertices $o \in O$ and $l \in L$, we use F_o and F'_l to denote their flower in these decompositions, respectively.

Let $H = G[\mathcal{O} \cup \mathcal{L}]$ be the induced packing graph of $\mathcal{F} = \mathcal{O} \cup \mathcal{L}$. The set \mathcal{F} is a $(2, 1)$ -shallow cover of G , and [Lemma 35](#) implies that H has polynomial expansion of order $k + 2$. By [Corollary 24 \(A\)](#), H has $\lambda = O((1/\varepsilon)^{2k+6} \log^{2k+5}(1/\varepsilon))$ -division $\mathcal{W} = \{C_1, \dots, C_m\}$ with a set of boundary vertices B , and total excess $(\varepsilon/4) |\mathcal{F}| \leq (\varepsilon/4) (|\mathcal{O}| + |\mathcal{L}|) \leq (\varepsilon/2) |L|$. For $i = 1, \dots, m$, let

- (i) $O_i = \{o \in O \mid F_o \in \mathcal{O} \cap C_i\}$,
- (ii) $L_i = \{l \in L \mid F'_l \in (\mathcal{L} \cap C_i) \setminus B\}$, and
- (iii) $B_i = B \cap C_i$.

Fix i , and consider the set $L' = (L \setminus L_i) \cup O_i$. If a vertex $v \in R$ is not dominated by $L \setminus L_i$, then $v \in F'_l \subseteq N(l) \cup \{l\}$ for some $l \in L_i$, and $v \in F_o \subseteq N(o) \cup \{o\}$ for some $o \in O$ with F'_l adjacent to F_o in H . The cluster F'_l is an interior vertex of C_i , so F_o must be in the cluster C_i , and $o \in O_i$. Therefore, the alternative solution L' dominates v , and overall, L' dominates R .

Since L is λ -locally minimal, and the exchange size is $|L \Delta L'| = |L_i \cup O_i| \leq |C_i| \leq \lambda$, the new solution L' is at least as large as L . Expanding $|L| \leq |L'| = |(L \setminus L_i) \cup O_i| = |L| - |L_i| + |O_i|$, we have $|L_i| \leq |O_i|$. Summed over all the clusters W_i , we have,

$$|L| \leq \sum_{i=1}^m (|L_i| + |B_i|) \leq \sum_{i=1}^m (|O_i| + |B_i|) \leq |O| + 2 \text{excess}(\mathcal{W}) \leq |O| + \frac{\varepsilon}{2} |L|.$$

Solving for $|L|$, we conclude that $|L| \leq |O| / (1 - \varepsilon/2) \leq (1 + \varepsilon) |O|$, as desired. \blacksquare

3.3.5. Extensions – multi-cover and reach

One can naturally extend dominating set in the following ways:

- (A) **Demands:** For every $v \in R$, there is an integer $\delta(v) \geq 0$, which is the *demand* of v ; that is, v has to be adjacent to at least $\delta(v)$ vertices in the dominating set. In the context of set cover, this is known as the multi-cover problem, see [CCH12]. Let $\hat{\delta} = \max_{v \in R} \delta(v)$ be the *demand* of the given instance.
- (B) **Reach:** Instead of the dominating set being adjacent to the vertices that are being covered, for every vertex $v \in R$ one can associate a distance $\tau(v) \geq 1$ – which is the maximum number of hops the dominating vertex can be away from v in the given graph. The *reach* of the given instance is $\hat{\tau} = \max_{v \in R} \tau(v)$.

Thus, a vertex v with demand $\delta(v)$ and reach $\tau(v)$, requires that any dominating set would have $\delta(v)$ vertices in edge distance at most $\tau(v)$ from it.

Lemma 41. *Let $G = (V, E)$ be a graph with n vertices and with polynomial expansion of order k , sets $R \subseteq V$ and $D \subseteq V$, such that D dominates R , and let $\varepsilon > 0$ be fixed. Furthermore, assume that for each vertex $v \in R$, there are associated demand and reach, where the reach $\hat{\tau}$ and demand $\hat{\delta}$ of the given instances are bounded by a constant.*

Then, for $\lambda = O(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon))$, the λ -local search algorithm computes, in $n^{O(\lambda)}$ time, a $(1 + \varepsilon)$ -approximation for the smallest cardinality subset of D that dominates R under the reach and demand constraints.

Proof: Let \prec be an arbitrary ordering on the vertices of G . For a set of vertices $X \subseteq V$ and a vertex $z \in V$, let $\text{NN}_k(z, X)$ be the k closest vertices to z in X , with respect to the length of the shortest path in G , and resolving ties by \prec . The ordering \prec ensures that $\text{NN}_k(z, X)$ is uniquely defined for any vertex in the graph.

In the following argument, fix a set $X \subseteq D$ that dominates R and complies with the given constraints, and assign every vertex of $u \in R$ to each of the vertices of $\text{NN}_{\delta(u)}(u, X)$. For a vertex $v \in X$, let $S(v)$ be the set of vertices assigned to it. For each vertex $v \in X$, let T_v be the minimal subtree of the BFS tree rooted at v that includes all the vertices of $S(v) \cup \{v\}$. The *flower* $C_v = V(T_v)$ is $\hat{\tau}$ -shallow in G . Let $\mathcal{F} = \mathfrak{F}(X, R) = \{C_v \mid v \in X\}$ be the resulting *flower decomposition* of X . (Note that a vertex $v \in V \setminus R$ with 0 covering demand may still be assigned to several flowers.)

We claim that a vertex z of G is covered at most $\hat{\delta}$ times by the flowers of \mathcal{F} . More precisely, we prove that z is covered by a flower C_v only if $v \in \text{NN}_{\hat{\delta}}(z, X)$. For the sake of contradiction, suppose

$z \in C_v$ and that $v \notin \text{NN}_{\hat{\delta}}(z, X)$. Then z is not assigned to v , so there must be a vertex u assigned to z and an associated shortest path $\pi_{uv} = \pi_{uz} \mid \pi_{zv}$ from u to v through z , where π_{uz} is the subpath from u to z and π_{zv} is the subpath from z to v . Since $v \in \text{NN}_{\hat{\delta}}(u, X) \setminus \text{NN}_{\hat{\delta}}(z, X)$, and both sets $\text{NN}_{\hat{\delta}}(u, X)$ and $\text{NN}_{\hat{\delta}}(z, X)$ have the same cardinality, there exists another vertex $v' \in \text{NN}_{\hat{\delta}}(z, X) \setminus \text{NN}_{\hat{\delta}}(u, X)$. Let $\sigma_{zv'}$ be the shortest-path from z to v' . By construction of $\text{NN}_{\hat{\delta}}(z, X)$, either $\|\sigma_{zv'}\| < \|\pi_{zv}\|$, or $\|\sigma_{zv'}\| = \|\pi_{zv}\|$ and $v' \prec v$. This implies that either $\|\pi_{uz} \mid \sigma_{zv'}\| < \|\pi_{uz} \mid \pi_{zv}\|$, or $v' \prec v$ and $\|\pi_{uz} \mid \sigma_{zv'}\| = \|\pi_{uz} \mid \pi_{zv}\|$, where \mid denotes concatenation of paths. In any case, if ties are broken by \prec , then v' is closer to u than v is, a contradiction to the premise that $v \in \text{NN}_{\hat{\delta}}(u, X)$ and $v' \notin \text{NN}_{\hat{\delta}}(u, X)$. Thus, if z is in a flower C_v , then $v \in \text{NN}_{\hat{\delta}}(z, X)$.

Now, consider the local solution L and the optimal solution O . Let $\mathcal{O} = \text{⊛}(O, R)$ and $\mathcal{L} = \text{⊛}(L, R)$ be the flower decompositions of the local and optimal solutions, respectively. Each flower decomposition includes an element at most $\hat{\delta}$ times, so the combined collection $\mathcal{F} = \mathcal{O} \cup \mathcal{L}$ is a $(2\hat{\delta}, \hat{\tau})$ -shallow packing. By [Lemma 35](#), the induced packing graph $H = G[\mathcal{F}]$ has polynomial expansion of order $k + 2$. We now follow the argument used in the proof of [Theorem 40](#), providing the details for the sake of completeness.

Let $\lambda = O(\varepsilon^{-2k+6} \log^{2k+5}(1/\varepsilon))$. There is a λ -division of H into clusters $\mathcal{C}_1, \dots, \mathcal{C}_m \subseteq \mathcal{F}$, with $\mathcal{B} \subseteq \mathcal{F}$ boundary vertices and total excess $|\mathcal{B}| \leq (\varepsilon/4) |\mathcal{F}|$. For $i = 1, \dots, m$, let

- (i) $\mathcal{O}_i = \mathcal{O} \cap \mathcal{C}_i$,
- (ii) $\mathcal{L}_i = (\mathcal{L} \cap \mathcal{C}_i) \setminus \mathcal{B}$, and
- (iii) $\mathcal{B}_i = \mathcal{B} \cap \mathcal{C}_i$.

Fix i , and consider the cover $\mathcal{L}' = (\mathcal{L} \setminus \mathcal{L}_i) \cup \mathcal{O}_i$. Consider a vertex $v \in V$ covered by a flower in $\mathcal{L} \setminus \mathcal{L}'$ (i.e., the vertex ‘‘lost’’ coverage in this potential exchange). Then v is covered by a flower $F \in \mathcal{L}_i$; that is, by a flower that corresponds to a vertex of H that is internal to \mathcal{C}_i . Any flower $F' \in \mathcal{F}$ that covers v is adjacent to F in H , by the definition of H and as F and F' share a vertex. As F is internal to \mathcal{C}_i , all the flowers of \mathcal{F} that cover v are in \mathcal{C}_i , and in particular, all the flowers covering v in the optimal solution belong to \mathcal{O}_i . Thus, the coverage provided by \mathcal{L}' meets the demand and reach requirements of v . The rest of the argument now follows the proof of [Theorem 40](#). \blacksquare

3.3.6. Extension: Connected dominating set

The algorithms of [Theorem 40](#) and [Lemma 41](#) can be extended to handle the additional constraint that the computed dominating set is also connected. In this setting, the local search algorithm only considers beneficial exchanges that result in a connected dominating set.

Lemma 42. *Let $G = (V, E)$ be a graph with n vertices and polynomial expansion of order k , and let $D \subseteq V$ be a connected dominating set. For each vertex $v \in V$, let $\delta(v) \geq 1$ be its associated demands, and let $\hat{\delta} = \max_{v \in V} \delta(v)$ be bounded by a constant (here, the dominating set has to dominate all the vertices in the graph). Then, for $\lambda = O(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon))$, the λ -local search algorithm computes, in $n^{O(\lambda)}$ time, a $(1 + \varepsilon)$ -approximation for the smallest cardinality subset of D that is connected and dominates V under the demand constraints.*

Proof: We extend the notations and argument used in [Theorem 40](#). To recap, let $O \subseteq D$ and $L \subseteq D$ be the optimal and locally minimum sets dominating R , respectively. Let $\mathcal{O} = \text{⊛}(O, R)$ and $\mathcal{L} = \text{⊛}(L, R)$ be the corresponding flower decompositions (see [Definition 39](#)). In the following, for vertices $o \in O$ and $l \in L$, we use F_o and F_l' to denote their flower in these decompositions, respectively.

Let $H = G[\mathcal{O} \cup \mathcal{L}]$ be the induced packing graph of $\mathcal{F} = \mathcal{O} \cup \mathcal{L}$. As before, we can apply [Corollary 24 \(A\)](#) to H to generate a $\lambda = O((1/\varepsilon)^{2k+6} \log^{2k+5}(1/\varepsilon))$ -division $\mathcal{W} = \{C_1, \dots, C_m\}$ with a set of boundary vertices B , and total excess $(\varepsilon/4) |\mathcal{F}| \leq (\varepsilon/4) (|\mathcal{O}| + |\mathcal{L}|) \leq (\varepsilon/2) |L|$. For $i = 1, \dots, m$, let

- (i) $O_i = \{o \in O \mid F_o \in \mathcal{O} \cap C_i\}$,
- (ii) $L_i = \{l \in L \mid F'_l \in (\mathcal{L} \cap C_i) \setminus B\}$, and
- (iii) $B_i = B \cap C_i$.

Fix i , and consider the set $L' = (L \setminus L_i) \cup O_i$. By the exact same argument as [Theorem 40](#), L' is a dominating set. However, L' may not necessarily be connected.

Let $B_i = \text{heads}(\mathcal{B}_i)$ be the set of head vertices of the boundary flowers of the i th cluster. Because the removed patch L_i is only connected to the rest of L via the boundary vertices B_i , each component of L contains at least one boundary vertex in $B_i \cap L$. Similarly, each component of O_i contains at least one boundary vertex in B_i . Together, every component of L' contains at least one vertex in B_i , so L' has at most $|B_i| \leq \lambda$ components.

Consider the shortest path π_{xy} within D between any two vertices $x, y \in L'$ that are in separate components of L' . By minimality of π_{xy} , the interior vertices of π_{xy} are not in L' . If π_{xy} has more than 4 vertices, then there exists an intermediate vertex $v \in \pi_{xy}$ that is adjacent to neither x nor y . Write $\pi_{xy} = \pi_{xv} \cup \pi_{vy}$, where π_{xv} is the subpath from x to v and π_{vy} is the subpath from v to y . Both subpaths π_{xv} and π_{vy} contain at least two edges. Since $\delta(v) \geq 1$, v is adjacent to some vertex $z \in L'$. Since x and y lie in different components, z lies in a different component from either x or y . If x and z lie in different components, then the path consisting of π_{xv} followed by the edge from v to z is a shorter path than π_{xy} , a contradiction. A similar contradiction arises if z and y lie in different components. It follows, by contradiction, that π_{xy} has at most 4 vertices, all of which lie in D . By adding the entire path π_{xy} to L' , we can connect these two components by adding at most 2 vertices from D .

By repeatedly connecting the closest pair of components of L' like that, we can augment L' to a connected dominating set L'' while adding at most $2|B_i| \leq 2\lambda$ vertices. If we expand our search size to $\lambda' = 3\lambda$, then L'' is a connected dominating set with $|L'' \Delta L| \leq \lambda'$, and the local optimality of L implies that $|L_i| \leq |O_i| + 2|B_i|$. As in the previous proofs, summing this inequality over all i implies the claim. ■

[Lemma 42](#) extends to constantly bounded reach with an added assumption.

Lemma 43. *Let $G = (V, E)$ be a graph with n vertices and with polynomial expansion of order k , and let $D \subseteq V$ be a given set. Assume that*

- (i) *for each vertex $v \in V$, there are associated demand $\delta(v) \geq 1$ and reach $\tau(v)$ constraints,*
- (ii) *$\hat{\delta} = \max_v \delta(v) = O(1)$ and $\hat{\tau} = \max_v \tau(v) = O(1)$,*
- (iii) *the set D is a valid dominating set complying with the demand and reach constraints,*
- (iv) *for any two vertices $u, v \in D$, the shortest path (in the number of edges) in G between u and v is contained in $G|_D$.*

Then, for $\lambda = O(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon))$, the λ -local search algorithm computes, in $n^{O(\lambda)}$ time, a $(1 + \varepsilon)$ -approximation for the smallest cardinality subset of D that is connected and dominates V under the reach and demand constraints.

Proof: The same proof as that of [Lemma 42](#) goes through, except now the shortest paths between distinct components can be shown to have length at most $2(\hat{\tau} + 1)$ vertices. Condition (iv) is necessary to keep these paths lying in D . The search size is increased by a factor of $2\hat{\tau}$ instead of 2, which is only a constant factor difference. ■

3.3.7. Discussion

Observation 44 (PTAS for vertex cover for graphs with polynomial expansion). *The algorithm of [Theorem 40](#) can be used to get a PTAS for vertex cover. Indeed, let $G = (V, E)$ be an undirected graph*

with polynomial expansion. We introduce a new vertex in the middle of every edge of G , and let H be the resulting graph, with R be the set of new vertices. Clearly, replacing an edge by a path of length two changes the expansion of a graph only slightly, see [Definition 14](#), and in particular, H has polynomial expansion. Now, solving the dominating subset for R as the set required covering, and V as the initial dominating set, in the graph H solves the original vertex cover problem in the original graph. The desired PTAS now follows from [Theorem 40](#).

Remark 45 (PTAS for graphs with subexponential expansion). As noted in [Remark 23](#), one can still obtain $f(1/\epsilon)$ -divisions for some (larger) function f in graphs with hereditary separators size $O(n/\log^{O(1)} n)$. To this end, one can verify (by the same proof as [Theorem 18](#), see [[NO08b](#), [HQ16](#)]) that for a small constant c , if a graph class \mathcal{C} has expansion $\varphi(t) = O(\exp(c' \cdot t^{c''}))$, for c' and c'' sufficiently small constants, then \mathcal{C} has separators of the desired size $O(n/\log^{O(1)} n)$. Thus, the above approximation algorithms yield a PTAS (with much worse dependence on ϵ) for any graph class \mathcal{C} with subexponential expansion $d_t(\mathcal{C}) = O(\exp(c' \cdot t^{c''}))$, where c' and c'' are some constants. We are not aware of any natural graphs in this class that do not have polynomial expansion.

3.4. Geometric applications

The above implies PTAS's for dominating set type problems on low-density graphs. Let \mathcal{U} be a low-density collection of objects in \mathbb{R}^d and P a collection of points. Two natural geometric optimization problems in this setting are:

- (A) **Discrete hitting set**: Compute the minimum cardinality set $Q \subseteq P$ such that for every $f \in \mathcal{U}$, we have $Q \cap f \neq \emptyset$. That is, every object of \mathcal{U} is stabbed by some point of Q .
If we consider the natural intersection graph $G = G_{P\cup\mathcal{U}}$ and the sets $D = P$ and $R = \mathcal{U}$, then this is an instance of dominating subset problem. The algorithm of [Theorem 40](#) applies because G is low density and therefore has polynomial expansion.
- (B) **Discrete set cover**: Compute the smallest cardinality set $\mathcal{V} \subseteq \mathcal{U}$ such that for every point $p \in P$, we have $p \in \bigcup_{f \in \mathcal{V}} f$. That is, all the points of P are covered by objects in \mathcal{V} . Setting $D = \mathcal{U}$ and $R = P$ (i.e., flipping the sets in the hitting set case), and arguing as above, implies a PTAS.

For these geometric optimization problems, we can improve the running time of [Theorem 40](#) by applying the stronger separator theorem for low-density graphs.

Theorem 46. *Let \mathcal{U} be a collection of m objects in \mathbb{R}^d with density ρ , P be a set of n points in \mathbb{R}^d , and let $\epsilon > 0$ be a parameter. Then, for $\lambda = O(\rho/\epsilon^d)$, the local search algorithm, with exchanges of size λ implies the following:*

- (A) *An approximation algorithm that, in $O(mn^{O(\lambda)})$ time, computes a set $Q \subseteq P$ that is an $(1 + \epsilon)$ -approximation for the smallest cardinality set that hits \mathcal{U} .*
- (B) *An approximation algorithm that, in $O(nm^{O(\lambda)})$ time, computes a set $\mathcal{V} \subseteq \mathcal{U}$ that is an $(1 + \epsilon)$ -approximation for the smallest cardinality set that covers P .*

Proof: Since points have zero diameter, the union $\mathcal{U} \cup P$ also has density $\rho + 1$. This reduces geometric hitting set and discrete geometric set cover to dominating subset problem on the intersection graph of $G = G_{\mathcal{U} \cup P}$.

The approximation algorithm is described in [Theorem 40](#) (applied to G). Here we can do slightly better, using smaller exchange size, as the graph G has low density. To this end, observe that the analysis of [Theorem 40](#) argues about the induced packing graph of G for some $(2, 1)$ -shallow packing \mathcal{G} . By [Lemma 38](#), the graph $H = G[\mathcal{G}]$ has density $O(\rho 2^d) = O(\rho)$. Thus, by [Corollary 24](#) (B), H has a

λ -division with excess $(\varepsilon/4) |V(H)|$, where $\lambda = O(\rho/\varepsilon^d)$. The algorithm of [Theorem 40](#) modified to use these improved divisions implies the result. ■

Remark 47. To our knowledge, the algorithms of [Theorem 46](#) are the first PTAS's for discrete hitting set and discrete set cover with shallow fat triangles and similar fat objects. Previously, such algorithms were known only for disks and points in the plane.

Remark 48. [Lemma 43](#) leads to PTAS for connected dominating sets for low density objects. Such problems were studied for unit disk graphs and disk graphs in two dimensions.

4. Hardness of approximation

Some of the results of this section appeared in an unpublished manuscript [[Har09](#)]. Chan and Grant [[CG14b](#)] also prove some related hardness results, which were (to some extent) a followup work to the aforementioned manuscript.

4.1. A review of complexity terms

The *exponential time hypothesis* (*ETH*) [[IP01](#), [IPZ01](#)] is that 3SAT can not be solved in time better than $2^{\Omega(n)}$, where n is the number of variables. The *strong exponential time hypothesis* (*SETH*), is that the time to solve k SAT is at least $2^{c_k n}$, where c_k converges to 1 as k increases.

A problem that is APX-HARD does not have a PTAS unless $P = NP$. For example, it is known that Vertex Cover is APX-HARD even for a graph with maximum degree 3 [[ACG+99](#)]. Thus, showing that a problem is APX-HARD implies that one can not do better than a constant approximation. Specifically, if one can get a $(1 + \varepsilon)$ -approximation for such a problem, for any constant $\varepsilon > 0$, then one can $(1 + \varepsilon)$ -approximate 3SAT (for the max version of 3SAT, the purpose is to maximize the number of clauses satisfied). By the PCP Theorem, this would imply an exact algorithm for 3SAT.

Observation 49. *Consider an instance of 3SAT with $m = c \log^2 n$ variables, for some constant $c > 0$ sufficiently large. ETH implies that we cannot solve this instance in polynomial time (in n), since the running time required to solve this instance is $2^{\Omega(m)}$, which is super polynomial. (This argument works for any function $f(n)$ such that $\log n = o(f(n))$.)*

This innocuous observation has a surprising implication – we cannot even $(1 - \varepsilon)$ approximate a solution for such an instance by the PCP result. Namely, ETH implies that even polylogarithmic sized instances cannot be solved in time polynomial in n .

Observation 50. *Showing that a problem X is APX-HARD implies that:*

- (A) *The problem X does not have a PTAS (unless $P = NP$).*
- (B) *Under ETH, the problem X does not have a QPTAS, where a QPTAS is an $(1 + \varepsilon)$ -approximation algorithm with running time $n^{O(\text{poly}(\log n, 1/\varepsilon))}$.*
- (C) *Furthermore, under ETH, instances of X of size $\log^{\Omega(1)} n$ cannot be approximated to within a $(1 \pm \varepsilon)$ -multiplicative factor in polynomial time in n .*

4.2. Discrete hitting set for fat triangles

In the *fat-triangles discrete hitting set problem*, we are given a set of points in the plane P and a set of fat triangles \mathcal{T} , and want to find the smallest subset of P such that each triangle in \mathcal{T} contains at least one point in the set.

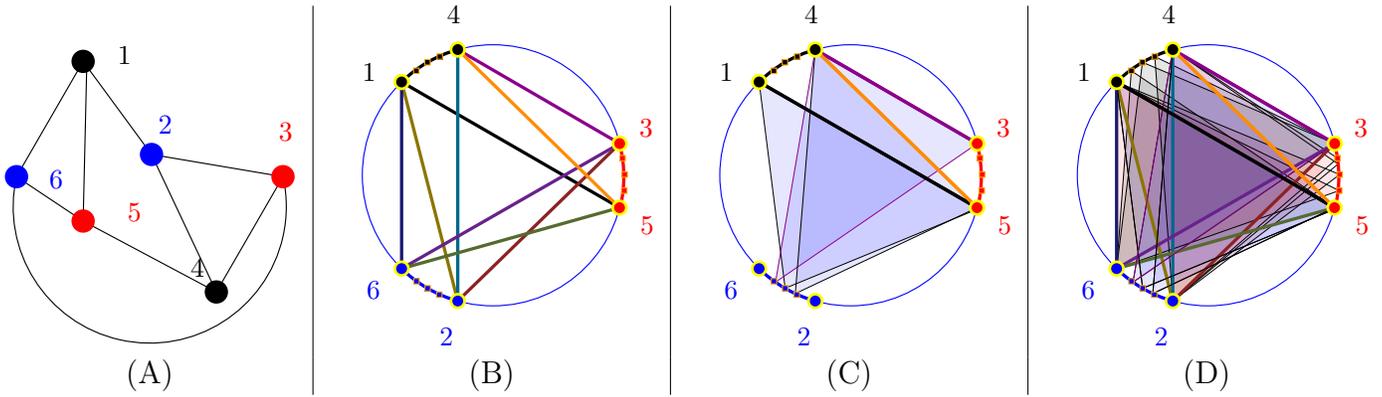


Figure 4.1: Illustration of the proof of [Lemma 51](#): (A) A 3-regular graph with its 3 coloring. (B) Placing the vertices on a circle. (C) Three edges and their associated triangles. (D) All the triangles.

Lemma 51. *There is no PTAS for the fat-triangle discrete hitting set problem, unless $P = NP$. One can prespecify an arbitrary constant $\delta > 0$, and the claim would hold true even if the following conditions hold on the given instance (P, \mathcal{T}) :*

- (A) *Every angle of every triangle in \mathcal{T} is between $60 - \delta$ and $60 + \delta$ degrees.*
- (B) *No point of P is covered by more than three triangles of \mathcal{T} .*
- (C) *The points of P are in convex position.*
- (D) *All the triangles of \mathcal{T} are of similar size. Specifically, each triangle has side length in the range (say) $(\sqrt{3} - \delta, \sqrt{3} + \delta)$.*
- (E) *The points of P are a subset of the vertices of the triangles of \mathcal{T} .*

Proof: Let $G = (V, E)$ be a connected instance of **Vertex Cover** which has maximum degree three, and it is not an odd cycle. We remind the reader that **Vertex Cover** is **APX-HARD** for such instances [[ACG+99](#)].

By Brook's theorem [[CR15](#)]⁵, this graph is three colorable, and let V_1, V_2, V_3 be the partition of V by their colors. Let p_1, p_2, p_3 be three points on the unit circle that form a regular triangle. For $i = 1, 2, 3$, place a circular interval J_i centered at p_i of length $\delta/100$. Now, for $i = 1, 2, 3$, we place the vertices of V_i as distinct points in J_i .

Let $Q_0 = V$ and $m = |E(G)|$. For $i = 1, \dots, m$, let $u_i v_i$ be the i th edge of G . Assume, for the sake of simplicity of exposition, that $u_i \in V_1$ and $v_i \in V_2$. Pick an arbitrary point q_i in $J_3 \setminus Q_{i-1}$, and create the triangle $T_i = \Delta u_i v_i q_i$. Set $Q_i = Q_{i-1} \cup \{q_i\}$, and continue to the next edge.

At the end of this process, we have m triangles $\mathcal{T} = \{T_1, \dots, T_m\}$ that are arbitrarily close to being regular triangles, and all their edges are arbitrarily close to being of the same length, see [Figure 4.1](#). It is easy to verify that a minimum cardinality set of points $U \subseteq V$ that hits all the triangles in \mathcal{T} is a minimum vertex cover of G . ■

4.3. Friendly geometric set cover

Let P be a set of n points in the plane, and \mathcal{F} be a set of m regions in the plane, such that

- (I) the shapes of \mathcal{F} are convex, fat, and of similar size,
- (II) the boundaries of any pair of shapes of \mathcal{F} intersect in at most 6 points,
- (III) the union complexity of any m shapes of \mathcal{F} is $O(m)$, and

⁵Brook's theorem states that any connected undirected graph G with maximum degree Δ , the chromatic number of G is at most Δ unless G is a complete graph or an odd cycle, in which case the chromatic number is $\Delta + 1$.

(IV) any point of P is covered by a constant number of shapes of \mathcal{F} .

We are interested in finding the minimum number of shapes of \mathcal{F} that covers all the points of P . This variant is the *friendly geometric set cover* problem, which is a variant of the following classical problem.

Definition 52. In the *Set Cover* problem, an instance is a ground set U , and a family \mathcal{F} of subsets of U , and the task is to find a minimum cardinality collection of sets of \mathcal{F} , such that their union is equal to U .

We need the following (known and easy) helper lemma.

Lemma 53. Consider an instance (U, \mathcal{F}) of *Set Cover* with the following properties:

- (A) All sets of \mathcal{F} are of size at most three.
- (B) Every element of U participates in exactly two sets of \mathcal{F} .
- (C) The elements of U are colored by four colors, such that every set of \mathcal{F} has no two elements of the same color.

Such an instance is *APX-HARD*.

Proof: Consider an instance of *Vertex Cover* on a graph G with maximum degree 3. Such instances are *APX-HARD* [ACG+99], and can be reduced to the claimed variant of *Set Cover* as follows. Let the ground set U consist of the edges of G , and let every vertex $v \in V(G)$ correspond to its set of edges, $S_v = \{uv \in E(G) \mid u \in V(G)\}$. Since G has maximum degree 3, each set S_v has size at most 3. Any cover C of size t for the set system $\mathcal{X} = (U, \{S_v \mid v \in V(G)\})$, corresponds to a vertex cover of G of the same size. Moreover, every ground element in the set cover instance participates in exactly two sets – the sets associated with the two endpoints of the underlying edge.

As for the last property, the graph G has maximum degree 3, and by Vizing’s theorem [BM76], it is 4 edge-colorable⁶. With respect to the *Set Cover* problem, the ground set of the set system \mathcal{X} can be colored by 4 colors such that no set in the set system has a color appearing more than once. ■

Lemma 54. There is no *PTAS* for the friendly geometric set cover problem, unless $P = NP$.

Proof: Consider an instance of *Set Cover* per Lemma 53, consisting of a set $U = \{u_1, \dots, u_n\}$ of n elements, and a family $\mathcal{F} = \{S_1, \dots, S_m\}$ of m subsets of U , each of size at most three. Every element of U is covered by at most two sets of \mathcal{F} . We reduce (U, \mathcal{F}) into an instance of the friendly geometric set cover problem as follows.

We place n points equally spaced on the unit radius circle centered at the origin, and let $P = \{p_1, \dots, p_n\}$ be the resulting set of points. For each point $u_i \in U$, let $f(u_i) = p_i$. For each set $S_i \in \mathcal{F}$ (of size at most 3), we define the region

$$g_i = \mathcal{CH} \left(\text{disk} \left(1 - \frac{i}{10n^2m} \right) \cup f(S_i) \right),$$

where \mathcal{CH} is the convex hull, $f(S_i) = \cup_{x \in S_i} \{f(x)\}$, and $\text{disk}(r)$ denotes the disk of radius r centered at the origin. Visually, g_i is a disk with at most three teeth coming out of it, see Figure 4.2. Note that the boundary of two such shapes intersects in at most 6 points.

⁶Vizing’s theorem states that a graph with maximum degree Δ can be edge colored by $\Delta + 1$ colors. In this specific case, one can reach the same conclusion directly from Brook’s theorem. Indeed, in our case, the adjacency graph of the edges has degree at most 4, and it does not contain a clique of size 4. As such, this graph is 4-colorable, implying the original graph edges are 4-colorable.

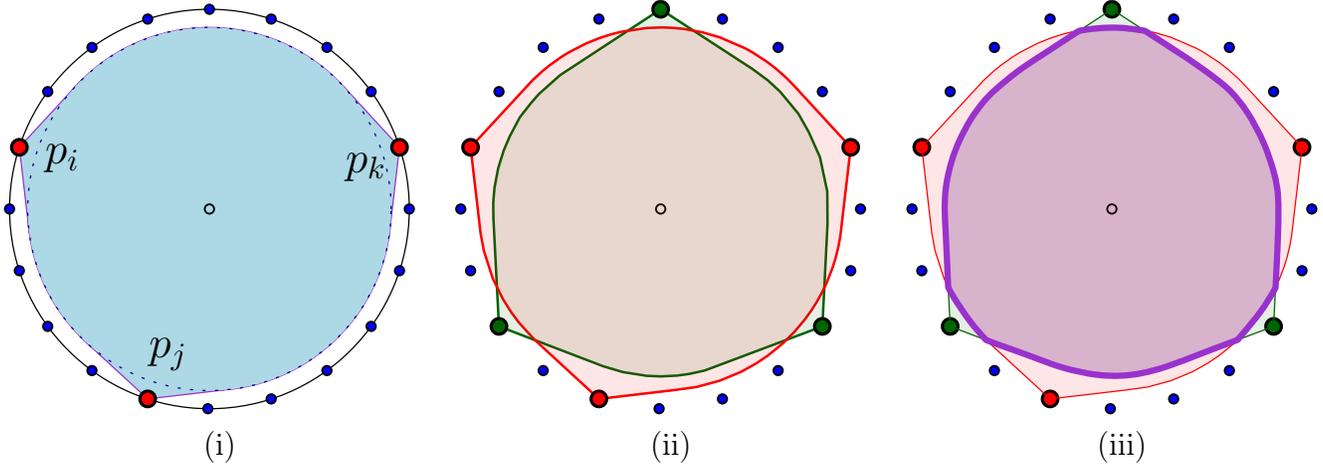


Figure 4.2: (i) A region g constructed for the set $S_t = \{u_i, u_j, u_k\}$. Observe that in the construction, the inner disk is even bigger. As such, no two points are connected by an edge of the convex-hull when we add in the inner disk to the convex-hull. As such, each point “contribution” to the region g is separated from the contribution of other points. (ii) How two such regions together. (iii) Their intersection.

Note, that every point of P is covered by exactly two such shapes, thus implying property (IV). It is now easy to verify that the resulting instance of geometric set cover $(P, \{g_1, \dots, g_m\})$ is friendly, and clearly any cover of P by these shapes can be interpreted as a cover of U by the corresponding sets of \mathcal{F} . Thus, a PTAS for the friendly geometric set cover problem implies a PTAS for the minimum k -set cover, which is impossible unless $P = NP$. ■

Remark. The computation model used above is the *real RAM* – which enables us to use irrational numbers for coordinates of points. However, all the reductions described here can be modified to use rational numbers by using a convex polygon that approximates the circle.

4.4. Set cover by fat triangles

In the *fat-triangle set cover problem*, specified by a set of points in the plane P and a set of fat triangles \mathcal{T} , one wants to find the minimum subset of \mathcal{T} such that its union covers all the points of P .

Lemma 55. *There is no PTAS for the fat-triangle set cover problem, unless $P = NP$. Furthermore, one can prespecify an arbitrary constant $\delta > 0$, and the claim would hold true even if the following conditions hold on the given instance (P, \mathcal{T}) :*

- (A) *The minimum angle of all the triangles of \mathcal{T} is larger than $45 - \delta$ degrees.*
- (B) *No point of P is covered by more than two triangles of \mathcal{T} .*
- (C) *The points of P are in convex position.*
- (D) *All the triangles of \mathcal{T} are of similar size. Specifically, each triangle has diameter in the range (say) $(2 - \delta, 2]$.*
- (E) *Each triangle of \mathcal{T} has two angles in the range $(45 - \delta, 45 + \delta)$, and one angle in the range $(90 - \delta, 90 + \delta)$.*
- (F) *The vertices of the triangles of \mathcal{T} are the points of P .*

Proof: Consider an instance (U, \mathcal{F}) of Set Cover as specified in Lemma 53. Let U_1, \dots, U_4 be the partition of U by their color.

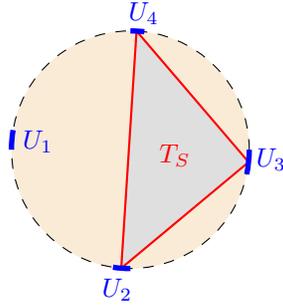


Figure 4.3

Let \mathcal{C} denote the circle of radius one centered at the origin. We place four relatively short arcs on \mathcal{C} , placed on the four intersection points of \mathcal{C} with the x and y axes, see Figure 4.3. Let I_1, \dots, I_4 denote these four circular intervals. We equally space the elements of U_i (as points) on the interval I_i , for $i = 1, \dots, 4$. Let P be the resulting set of points.

For every set $S \in \mathcal{F}_{\mathcal{X}}$, take the convex hull of the points corresponding to its elements as its representing triangle T_S . Note, that since the vertices of T_S lie on three intervals out of I_1, I_2, I_3, I_4 , it follows that it must be fat, for all $S \in \mathcal{F}_{\mathcal{X}}$. As such, the resulting set of triangles $\mathcal{T} = \{T_S \mid S \in \mathcal{F}_{\mathcal{X}}\}$ is fat, and clearly there is a cover of P by t triangles of \mathcal{T} if and only if the original set cover problem has a cover of size t .

Any triangle having its three vertices on three different intervals of I_1, \dots, I_4 is close to being an isosceles triangle with the middle angle being 90 degrees. As such, by choosing these intervals to be sufficiently short, any triangle of \mathcal{T} would have a minimum degree larger than, say, $45 - \delta$ degrees, and with diameter in the range between $2 - \delta$ and 2.

This is clearly an instance of the fat-triangle set cover problem. Solving it is equivalent to solving the original Vertex Cover problem, but since it is APX-HARD, it follows that the fat-triangle set cover problem is APX-HARD. ■

Remark 56. For fat triangles of similar size a constant factor approximation algorithm is known [CV07]. Lemma 55 implies that one can do no better. Naturally, it might be possible to slightly improve the constant of approximation provided by the algorithm of Clarkson and Varadarajan [CV07]. However, for fat triangles of different sizes, only a \log^* approximation is known [ABES14]. It is natural to ask if this can be improved.

Remark 57. An interesting property of the instances generated by Lemma 54 and Lemma 55 (and in also some of the results below) is that while the set of shapes in the resulting geometric instance is dense, the points that the instance has to cover are shallow – specifically they have depth two. This is in contrast to the situation where the whole plane is shallow, and a PTAS exists (i.e., Theorem 46).

4.4.1. Extensions

Lemma 58. *Given a set of points P in the plane and a set of circles \mathcal{F} , finding the minimum number of circles of \mathcal{F} that covers P is APX-HARD; that is, there is no PTAS for this problem.*

Proof: Slightly perturb the point set used in the proof of Lemma 55, so that no four points of it are co-circular. Let P denote the resulting set of points. For every set $S \in \mathcal{F}_{\mathcal{X}}$, we now take the circle passing through the three corresponding points. Clearly, this results in a set of circles (that are almost identical, but yet all different), such that finding the minimum number of circles covering the set P is equivalent to solving the original problem. ■

Lemma 59. *Given a set of points Q in \mathbb{R}^3 and a set of planes \mathcal{F} , finding the minimum number of planes of \mathcal{F} that covers Q is APX-HARD; that is, there is no PTAS for this problem.*

Proof: Let P be the point set and \mathcal{F} be the set of circles constructed in the proof of Lemma 58, and map every point in it to three dimensions using the mapping $f : (x, y) \rightarrow (x, y, x^2 + y^2)$. This is a standard lifting map used in computing planar Delaunay triangulations via convex-hull in three dimensions, see [BCKO08]. Let $Q = f(P)$ be the resulting point set.

It is easy to verify that a circle of $c \in \mathcal{F}$ is mapped by f into a curve that lies on a plane. We will abuse notations slightly, and use $f(c)$ to denote this plane. Let $\mathcal{G} = f(\mathcal{F})$. Furthermore, for a circle $c \in \mathcal{F}$, we have that $f(c \cap P) = f(c) \cap Q$. Namely, solving the set cover problem (Q, \mathcal{G}) is equivalent to solving the original set cover instance (P, \mathcal{F}) . ■

The recent work of Mustafa *et al.* [MR10] gave a QPTAS for set cover of points by disks (i.e., circles with their interior), and for set cover of points by half-spaces in three dimensions. Thus, somewhat surprisingly, the “shelled” version of these problems are harder than the filled-in version.

4.5. Independent set of triangles in 3D

Given a set \mathcal{U} of n objects in \mathbb{R}^d (say, triangles in 3d), we are interested in computing a maximum number of objects that are *independent*; that is, no pair of objects in this set (i.e., independent set) intersects. This is the geometric realization of the *independent set* problem for the intersection graph induced by these objects.

Lemma 60. *There is no PTAS for the maximum independent set of triangles in \mathbb{R}^3 , unless $P = NP$.*

Proof: The problem Independent Set is APX-HARD even for graphs with maximum degree 3 [ACG+99]. Let $G = (V, E)$ be a given graph with maximum degree 3, where $V = \{v_1, \dots, v_n\}$. We will create a set of triangles, such that their intersection graph is G .

If one spreads n points p_1, \dots, p_n on the positive branch of the moment curve in \mathbb{R}^3 [Sei91, EK03], their Voronoi diagram is *neighborly*; that is, every Voronoi cell is a convex polytope that shares a non-empty two dimensional boundary face with each of the other cells of the diagram. Let C_i denote the cell of the point p_i in this Voronoi diagram, for $i = 1, \dots, n$.

Now, for every vertex $v_i \in V$, we form a set P_i of (at most) three points, as follows. If $v_i v_j \in E$, then we place a point p_{ij} on the common boundary of C_i and C_j , and we add this point to both P_i and P_j . After processing all the edges in E , each point set P_i has at most three points, as the maximum degree in G is three.

For $i = 1, \dots, n$, let f_i be the triangle formed by the convex-hull of P_i (if P_i has fewer than three points then the triangle is degenerate).

Let $\mathcal{T} = \{f_1, \dots, f_n\}$. Observe that the triangles of \mathcal{T} are disjoint except maybe in their common vertices, as their interior is contained inside the interior of C_i , and the cells C_1, \dots, C_n are interior disjoint. Clearly $f_i \cap f_j \neq \emptyset$ if and only if $v_i v_j \in E$. Thus, finding an independent set in G is equivalent to finding an independent set of triangles of the same size in \mathcal{T} . We conclude that the problem of finding maximum independent set of triangles is APX-HARD, and as such does not have a PTAS unless $P = NP$. ■

Implicit in the above proof is that any graph can be realized as the intersection graph of convex bodies in \mathbb{R}^3 (we were a bit more elaborate for the sake of completeness and since we needed slightly more structure). This is well known and can be traced to a result of Tietze from 1905 [Tie05].

4.6. Hardness of approximation with respect to depth

Using [Observation 49](#), we can derive hardness of approximation results even for “small” instances. Here, we reconsider the *geometric set cover problem*: Given a set of objects \mathcal{U} in \mathbb{R}^d , and a set of points P , we would like to find minimum cardinality subset of the objects in \mathcal{U} that covers the points of P .

Lemma 61. *Assuming the exponential time hypothesis (ETH) (see [Section 4.1](#)), consider a given set of fat triangles \mathcal{T} of density ρ , and a set of points P , such that $|P| + |\mathcal{T}| = O(n)$. We have the following:*

- (A) *If $\rho = \Omega(\log^c n)$ then one cannot $(1 + \varepsilon)$ -approximate the geometric set cover (or geometric hitting set) instance (P, \mathcal{T}) in polynomial time, where c is a sufficiently large constant.*
- (B) *There is an absolute constant c'' , such that no $(1 + \varepsilon)$ -approximation algorithm for the geometric set cover (or hitting set) instance (P, \mathcal{T}) has running time $n^{\rho^{c''}/\varepsilon^{O(1)}}$.*

Proof: (A) Suppose we had such a PTAS, and consider an instance I of 3SAT of size at least $c' \log^2 n$, where c' is a sufficiently large constant. ETH implies that any algorithm solving such an instance must have running time at least $n^{\Omega(\log n)}$. On the other hand, the instance I can be converted to a set cover instance of fat triangles with polylog n triangles/points and polylog n density, by [Lemma 55](#). As such, a PTAS in this case, would contradict ETH.

(B) Consider a constant c'' sufficient small. By part (A), an instance of 3SAT with $\log^2 n$ variables, can be converted into an instance of geometric set cover (of fat triangles) with depth $\log^{c_3} n$ (where c_3 is some constant). But then, if $c'' < 1/c_3$, this implies that this instance be solved in polynomial time, contradicting ETH.

The same conclusions holds for geometric hitting set, by using [Lemma 51](#). ■

5. Conclusions

In this paper, we studied the class of graphs arising out of low density objects in \mathbb{R}^d , and showed that they are subclass of graphs with polynomial expansion. We provided PTAS’s for independent set and dominating set problems (and some variants) for such graphs. This gives rise to PTAS for some generic variants of these problems. Coupled with hardness results, we characterize the complexity of geometric variants of set cover and hitting set as a function of depth (for example, for fat triangles).

At this point in time, it seems interesting to better understand low density graphs. In particular, how exactly do they relate to graphs of low genus, and whether one can develop efficient approximation algorithms and hardness of approximations to other problems for this family of graphs. (In particular, there is strong evidence that low genus graphs are low density graphs.) For example, as a concrete problem, can one get a PTAS for TSP for low-density graphs or polynomial expansion graphs?

Acknowledgments. The authors thank Mark de Berg for useful discussions related to the problems studied in this paper. In particular, he pointed out the improved bound for [Lemma 19](#). We also thank the anonymous referees for their detailed and useful comments. We are particularly grateful to the anonymous referee of a previous version of this paper who pointed out the connection of our work to graphs with bounded expansion.

References

- [ABES14] B. Aronov, M. de Berg, E. Ezra, and M. Sharir. *Improved bounds for the union of locally fat objects in the plane*. *SIAM J. Comput.*, 43(2): 543–572, 2014.

- [ACG+99] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and approximation*. Springer-Verlag, 1999.
- [AES10] B. Aronov, E. Ezra, and M. Sharir. *Small-size ε -nets for axis-parallel rectangles and boxes*. *SIAM J. Comput.*, 39(7): 3248–3282, 2010.
- [And70] E. Andreev. *On convex polyhedra in lobachevsky spaces*. *Sbornik: Mathematics*, 10: 413–440, 1970.
- [APS08] P. K. Agarwal, J. Pach, and M. Sharir. *State of the Union—of Geometric Objects*. *Surveys in Discrete and Computational Geometry Twenty Years Later*. Ed. by J. E. Goodman, J. Pach, and R. Pollack. Vol. 453. Contemporary Mathematics. Amer. Math. Soc., 2008, pp. 9–48.
- [AST90] N. Alon, P. D. Seymour, and R. Thomas. *A separator theorem for graphs with an excluded minor and its applications*. *Proc. 22nd Annu. ACM Sympos. Theory Comput. (STOC)*, 293–299, 1990.
- [AW13] A. Adamaszek and A. Wiese. *Approximation schemes for maximum weight independent set of rectangles*. *Proc. 54th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, 400–409, 2013.
- [AW14] A. Adamaszek and A. Wiese. *A QPTAS for maximum weight independent set of polygons with polylogarithmic many vertices*. *Proc. 25th ACM-SIAM Sympos. Discrete Algs. (SODA)*, 400–409, 2014.
- [Bak94] B. S. Baker. *Approximation algorithms for NP-complete problems on planar graphs*. *J. Assoc. Comput. Mach.*, 41: 153–180, 1994.
- [BCKO08] M. de Berg, O. Cheong, M. Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications*. 3rd. Springer-Verlag, 2008.
- [Ber08] M. de Berg. *Improved bounds on the union complexity of fat objects*. *Discrete Comput. Geom.*, 40(1): 127–140, 2008.
- [BKSV02] M. de Berg, M. J. Katz, A. v. Stappen, and J. Vleugels. *Realistic input models for geometric algorithms*. *Algorithmica*, 34(1): 81–97, 2002.
- [BM76] J. A. Bondy and U. S. R. Murty. *Graph Theory with Applications*. North-Holland, 1976.
- [CCH12] C. Chekuri, K. Clarkson, and S. Har-Peled. *On the set multi-cover problem in geometric settings*. *ACM Trans. Algo.*, 9(1): 9, 2012.
- [CG09] J. Chalopin and D. Gonçalves. *Every planar graph is the intersection graph of segments in the plane: extended abstract*. *Proc. 41st Annu. ACM Sympos. Theory Comput. (STOC)*, 631–638, 2009.
- [CG14a] S. Cabello and D. Gajser. *Simple PTAS’s for families of graphs excluding a minor*. *CoRR*, abs/1410.5778, 2014.
- [CG14b] T. M. Chan and E. Grant. *Exact algorithms and APX-hardness results for geometric packing and covering problems*. *Comput. Geom. Theory Appl.*, 47(2): 112–124, 2014.
- [CG15] S. Cabello and D. Gajser. *Simple PTAS’s for families of graphs excluding a minor*. *Disc. App. Math.*, 189: 41–48, 2015.
- [CH12] T. M. Chan and S. Har-Peled. *Approximation algorithms for maximum independent set of pseudo-disks*. *Discrete Comput. Geom.*, 48(2): 373–392, 2012.

- [Cha03] T. M. Chan. *Polynomial-time approximation schemes for packing and piercing fat objects*. *J. Algorithms*, 46(2): 178–189, 2003.
- [CKR01] G. Calinescu, H. Karloff, and Y. Rabani. *Approximation algorithms for the 0-extension problem*. *Proc. 12th ACM-SIAM Sympos. Discrete Algs. (SODA)*, 8–16, 2001.
- [CR15] D. W. Cranston and L. Rabern. *Brooks’ theorem and beyond*. *J. Graph Theo.*, 80(3): 199–225, 2015.
- [CV07] K. L. Clarkson and K. R. Varadarajan. *Improved approximation algorithms for geometric set cover*. *Discrete Comput. Geom.*, 37(1): 43–58, 2007.
- [DEG16] K. Dutta, E. Ezra, and A. Ghosh. *Two proofs for shallow packings*. *Discrete Comput. Geom.*, 56(4): 910–939, 2016.
- [DN15] Z. Dvořák and S. Norin. *Strongly sublinear separators and polynomial expansion*. *ArXiv e-prints*, 2015. arXiv: 1504.04821 [math.CO].
- [EHS15] D. Eppstein, S. Har-Peled, and A. Sidiropoulos. *Approximate Greedy Clustering and Distance Selection for Graph Metrics*. *ArXiv e-prints*, 2015. arXiv: 1507.01555 [cs.CG].
- [EK03] J. Erickson and S. Kim. *Arbitrarily large neighborly families of congruent symmetric convex 3-polytopes*. *Discrete Geometry: In Honor of W. Kuperberg’s 60th Birthday*. Ed. by A. Bezdek. Lecture Notes Pure Appl. Math. Marcel-Dekker, 2003, pp. 267–278. eprint: arXiv: math.CO/0106095.
- [Epp00] D. Eppstein. *Diameter and treewidth in minor-closed graph families*. *Algorithmica*, 27: 275–291, 2000.
- [FG88] T. Feder and D. H. Greene. *Optimal algorithms for approximate clustering*. *Proc. 20th Annu. ACM Sympos. Theory Comput. (STOC)*, 434–444, 1988.
- [Fre87] G. N. Frederickson. *Fast algorithms for shortest paths in planar graphs, with applications*. *SIAM J. Comput.*, 16(6): 1004–1022, 1987.
- [FRT04] J. Fakcharoenphol, S. Rao, and K. Talwar. *A tight bound on approximating arbitrary metrics by tree metrics*. *J. Comput. Sys. Sci.*, 69(3): 485–497, 2004.
- [GKS14] M. Grohe, S. Kreutzer, and S. Siebertz. *Deciding first-order properties of nowhere dense graphs*. *Proc. 46th Annu. ACM Sympos. Theory Comput. (STOC)*, 89–98, 2014.
- [Gro03] M. Grohe. *Local tree-width, excluded minors, and approximation algorithms*. *Combinatorica*, 23(4): 613–632, 2003.
- [GRRR16] S. Govindarajan, R. Raman, S. Ray, and A. B. Roy. *Packing and covering with non-piercing regions*. *Proc. 24th Annu. Euro. Sympos. Alg. (ESA)*, 47:1–47:17, 2016.
- [Har09] S. Har-Peled. *Being fat and friendly is not enough*. *CoRR*, 2009. arXiv: 0908.2369 [cs.CG].
- [Har13] S. Har-Peled. *A simple proof of the existence of a planar separator*. *ArXiv e-prints*, 2013. arXiv: 1105.0103 [cs.CG].
- [Har14] S. Har-Peled. *Quasi-polynomial time approximation scheme for sparse subsets of polygons*. *Proc. 30th Annu. Sympos. Comput. Geom. (SoCG)*, 120–129, 2014.
- [Hås99] J. Håstad. *Clique is hard to approximate within $n^{1-\epsilon}$* . *Acta Mathematica*, 182(1): 105–142, 1999.
- [HKRS97] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian. *Faster shortest-path algorithms for planar graphs*. *J. Comput. Sys. Sci.*, 55(1): 3–23, 1997.

- [HQ15] S. Har-Peled and K. Quanrud. *Approximation algorithms for polynomial-expansion and low-density graphs*. *Proc. 23rd Annu. Euro. Sympos. Alg. (ESA)*, vol. 9294. 717–728, 2015.
- [HQ16] S. Har-Peled and K. Quanrud. *Notes on approximation algorithms for polynomial-expansion and low-density graphs*. *ArXiv e-prints*, 2016. arXiv: 1603.03098 [cs.CG].
- [HR15] S. Har-Peled and B. Raichel. *Net and prune: a linear time algorithm for euclidean distance problems*. *J. Assoc. Comput. Mach.*, 62(6): 44:1–44:35, 2015.
- [HT74] J. E. Hopcroft and R. E. Tarjan. *Efficient planarity testing*. *J. Assoc. Comput. Mach.*, 21(4): 549–568, 1974.
- [IP01] R. Impagliazzo and R. Paturi. *On the complexity of k -SAT*. *J. Comput. Sys. Sci.*, 62(2): 367–375, 2001.
- [IPZ01] R. Impagliazzo, R. Paturi, and F. Zane. *Which problems have strongly exponential complexity?* *J. Comput. Sys. Sci.*, 63(4): 512–530, 2001.
- [Kar72] R. M. Karp. *Reducibility among combinatorial problems*. *Complexity of Computer Computations*, 85–103, 1972.
- [KMM93] R. Klein, K. Mehlhorn, and S. Meiser. *Randomized incremental construction of abstract Voronoi diagrams*. *Comput. Geom. Theory Appl.*, 3: 157–184, 1993.
- [Koe36] P. Koebe. *Kontaktprobleme der konformen Abbildung*. *Ber. Verh. Sächs. Akademie der Wissenschaften Leipzig, Math.-Phys. Klasse*, 88: 141–164, 1936.
- [LT79] R. J. Lipton and R. E. Tarjan. *A separator theorem for planar graphs*. *SIAM J. Appl. Math.*, 36: 177–189, 1979.
- [LT80] R. J. Lipton and R. E. Tarjan. *Applications of a planar separator theorem*. *SIAM J. Comput.*, 9(3): 615–627, 1980.
- [Mat14] J. Matoušek. *Near-optimal separators in string graphs*. *Combin., Prob. & Comput.*, 23(1): 135–139, 2014.
- [MR10] N. H. Mustafa and S. Ray. *Improved results on geometric hitting set problems*. *Discrete Comput. Geom.*, 44(4): 883–895, 2010.
- [MRR14a] N. H. Mustafa, R. Raman, and S. Ray. *QPTAS for geometric set-cover problems via optimal separators*. *ArXiv e-prints*, 2014. arXiv: 1403.0835 [cs.CG].
- [MRR14b] N. H. Mustafa, R. Raman, and S. Ray. *Settling the APX-hardness status for geometric set cover*. *Proc. 55th Annu. IEEE Sympos. Found. Comput. Sci. (FOCS)*, 541–550, 2014.
- [MS09] M. Mendel and C. Schwob. *Fast C - K - R partitions of sparse graphs*. *Chicago J. Theor. Comput. Sci.*, 2009, 2009.
- [MTTV97] G. L. Miller, S. H. Teng, W. P. Thurston, and S. A. Vavasis. *Separators for sphere-packings and nearest neighbor graphs*. *J. Assoc. Comput. Mach.*, 44(1): 1–29, 1997.
- [NO08a] J. Nešetřil and P. Ossona de Mendez. *Grad and classes with bounded expansion I. decompositions*. *Eur. J. Comb.*, 29(3): 760–776, 2008.
- [NO08b] J. Nešetřil and P. Ossona de Mendez. *Grad and classes with bounded expansion II. algorithmic aspects*. *Eur. J. Comb.*, 29(3): 777–791, 2008.
- [NO12] J. Nešetřil and P. Ossona de Mendez. *Sparsity – Graphs, Structures, and Algorithms*. Vol. 28. Alg. Combin. Springer, 2012.

- [PA95] J. Pach and P. K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, 1995.
- [RS97] R. Raz and S. Safra. *A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP*. *Proc. 29th Annu. ACM Sympos. Theory Comput.* (STOC), 475–484, 1997.
- [Sei91] R. Seidel. *Exact upper bounds for the number of faces in d-dimensional Voronoi diagrams*. *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*. Ed. by P. Gritzman and B. Sturmfels. Vol. 4. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. Amer. Math. Soc., 1991, pp. 517–530.
- [SOBV98] A. F. van der Stappen, M. H. Overmars, M. de Berg, and J. Vleugels. *Motion planning in environments with low obstacle density*. *Discrete Comput. Geom.*, 20(4): 561–587, 1998.
- [SS85] J. T. Schwartz and M. Sharir. *Efficient motion planning algorithms in environments of bounded local complexity*. Report 164. New York, NY: Dept. Comput. Sci., Courant Inst. Math. Sci., New York Univ., 1985.
- [Sta92] A. Stappen. *Motion Planning Amidst Fat Obstacles*. PhD thesis. Netherlands: Utrecht University, 1992.
- [SW98] W. D. Smith and N. C. Wormald. *Geometric separator theorems and applications*. *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), 232–243, 1998.
- [Tie05] H. Tietze. *Über das problem der nachbargebiete im raum*. *Monatshefte Math.*, 15: 211–216, 1905.
- [Ver05] J.-L. Verger-Gaugry. *Covering a ball with smaller equal balls in \mathbb{R}^n* . *Discrete Comput. Geom.*, 33(1): 143–155, 2005.