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

Sariel Har-Peled[†] Kent Quanrud[‡]

September 23, 2015

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 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 interesting 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

Many classical optimization problems are intractable to approximate, let alone solve. Motivated by the discrepancy between the worst-case analysis and real-world success of algorithms, more realistic models of input have been developed, alongside algorithms that take advantage of their properties. In this paper, we investigate approximability of some classical optimization problems (e.g., set cover and independent set, among others) for two closely-related families of graphs: Graphs with *polynomially-bounded expansion*, and intersection graphs of geometric objects with *low-density*.

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

*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].

[†]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]	Exact version NP-Hard [FG88]
Fat triangles of same size	$O(1)$ [CV07]	APX-Hard: Lemma 4.4.1$_{\text{p22}}$ I.e., no PTAS possible.
Fat objects in $\mathbb{R}^2$	$O(\log^* \text{opt})$ [AdBES14]	APX-Hard: L4.4.1
Objects $\subseteq \mathbb{R}^d$, $O(1)$ density E.g. fat objects, $O(1)$ depth.	PTAS: Theorem 3.4.1	Exact version NP-Hard [FG88]
Objects with polylog density	QPTAS: Theorem 3.4.1	No PTAS under ETH Lemma 4.6.1$_{\text{p25}}$
Objects with density ρ in $\mathbb{R}^d$	PTAS: Theorem 3.4.1 RT: $n^{O(\rho^{(d+1)/d}/\epsilon^d)}$.	No $(1+\varepsilon)$-approx with RT $n^{\text{poly}(\log \rho, 1/\varepsilon)}$ assuming ETH: L4.6.1

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 that has running time $n^{O(\text{poly}(\log n, 1/\varepsilon))}$.

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 $n^{1-\varepsilon}$, for any fixed $\varepsilon > 0$, unless $\text{P} = \text{NP}$ [Hås96].

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 $\text{P} = \text{NP}$ [RS97].

1.1.2. Graph classes

Density. Informally, a set of objects in $\mathbb{R}^d$ is *low-density* if no ball 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]). A closely related geometric property to density is *fatness*. 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].

Intersection graphs. A set $\mathcal{F}$ of objects in $\mathbb{R}^d$ induces an *intersection graph* $G_{\mathcal{F}}$ having $\mathcal{F}$ as its the 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* 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 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. The class of graphs with polynomial expansion 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 sparsity of a graph (the ratio of edges to vertices) is not necessarily sufficient for tractability. For example, a clique (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).

Bounded expansion graphs are *nowhere dense graphs* [NO12, Section 5.4]. Grohe, Kreutzer and Siebertz recently showed that first-order properties are fixed-parameter tractable for nowhere dense graphs [GKS14]. In this paper, we study graphs of bounded expansion [NO12, Section 5.5], which intuitively requires a graph to not only be sparse, but have shallow minors that are sparse as well.

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 a PTAS for independent set and (continuous) piercing set for fat objects [Cha03], a PTAS for piercing half-spaces and pseudo-disks [MR10], a QPTAS for maximum weighted independent sets of polygons [AW13, AW14, Har14], and a QPTAS for Set Cover by pseudodisks [MRR14a], among others. Lastly, Cabello and Gajser [CG14] 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 with hereditary separators.** 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 2.3.1 and Lemma 3.1.3.

(C) **PTAS for independent/packing when the output is sparse.** More surprisingly, one get a PTAS even if the subgraph induced on the union of two solutions has polynomial expansion.

Objects	Approx. Alg.	Hardness
Disks/pseudo-disks	PTAS [MR10]	Exact version NP-Hard via point-disk duality [FG88]
Fat triangles of similar size.	$O(\log\log \mathrm{opt})$ [AES10]	APX-Hard: Lemma 4.2.1$_{\mathrm{p20}}$
Objects with $O(1)$ density.	PTAS: Theorem 3.4.1$_{\mathrm{p19}}$	Exact ver. NP-Hard [FG88]
Objects polylog density.	QPTAS: Theorem 3.4.1	No PTAS under ETH Lemma 4.6.1 / L4.2.1
Objects with density ρ in $\mathbb{R}^d$	PTAS: Theorem 3.4.1 run time $n^{O(\rho^{(d+1)/d}/\epsilon^d)}$	No $(1+\varepsilon)$-approx with RT $n^{\mathrm{poly}(\log\rho, 1/\varepsilon)}$ assuming ETH: L4.6.1

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.

Thus, while the input may not be sparse, as long as the output is sparse, one can get an efficient approximation algorithms, see Theorem 3.2.1.

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 merges locally objects that are close together, see Lemma 2.1.12. More generally, if one consider a collection of t-shallow subgraphs (i.e., graphs with edge distance radius t) 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 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 3.3.10.

(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 3.3.13.

(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 our approximations 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.

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. Appendix A contains some proofs that are provided for the sake of completeness.

2. Preliminaries

2.1. Low-density graphs

Definition 2.1.1. 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 $G_{|X} = \bigl(X, \{uv \mid u, v \in X, \text{ and } uv \in E\}\bigr)$.

Definition 2.1.2. 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 an edge between two objects $f, g \in \mathcal{U}$ if they intersect; that is, formally $G_{\mathcal{U}} = \bigl(\mathcal{U}, \{fg \mid f, g \in \mathcal{U} \text{ and } f \cap g \neq \emptyset\}\bigr)$.

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

Definition 2.1.3. A set of objects $\mathcal{U}$ in $\mathbb{R}^d$ (not necessarily convex or connected) has ***density*** ρ if any ball b intersects at most ρ objects in $\mathcal{U}$ with diameter larger than the diameter of b. The minimum such quantity is denoted by density$(\mathcal{U})$. If ρ is a constant, then $\mathcal{U}$ has ***low density***.

Any 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***. The class of all graphs that are ρ-dense and are induced by objects in $\mathbb{R}^d$ is denoted by $\mathcal{C}^d_\rho$.

Definition 2.1.4. A graph G is k***-degenerate*** if any subgraph of G has a vertex of degree at most k.

Observation 2.1.5. *A ρ-dense graph is $(\rho-1)$-degenerate (with degree $\rho-1$ attained by the object with smallest diameter). Thus, a ρ-dense graph with n vertices has at most $(\rho-1)n$ edges.*

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, we have $\mathrm{vol}(b \cap g) \geq \alpha \, \mathrm{vol}(b)$ [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.

Lemma 2.1.6. *A set $\mathcal{F}$ of α-fat convex 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 bounded density.*

Proof: Let $b = \mathbb{b}(p, r)$ be any ball in $\mathbb{R}^d$, and consider an α-fat object $g \in \mathcal{U}$ that intersects b and has $\mathrm{diam}(g) > \mathrm{diam}(b) = 2r$.

①There are several different, but roughly equivalent, definitions of fatness in the literature, see de Berg [dB08] and the followup work by Aronov *et al.* [AdBES14] for some recent results. In particular, our definition here is what de Berg refers to as being *locally fat*.

A ball $\mathbb{b}(q, r)$ centered at a point q that is in $g \cap b$ does not contain g, as $\operatorname{diam}(\mathbb{b}(q,r)) < \operatorname{diam}(g)$. As such, by the definition of α-fatness, we have $\operatorname{vol}\bigl(\mathbb{b}(q,r)\bigr) \geq \operatorname{vol}\bigl(\mathbb{b}(q,r) \cap g\bigr) \geq \alpha \operatorname{vol}\bigl(\mathbb{b}(q,r)\bigr) = \alpha \operatorname{vol}(b)$. Furthermore, $\mathbb{b}(q,r)$ is contained in the ball $b' = \mathbb{b}(p, 2r)$, and as such

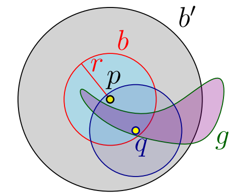

$$\operatorname{vol}\Bigl(b' \cap g\Bigr) \geq \operatorname{vol}\Bigl(\mathbb{b}(q,r) \cap g\Bigr) \geq \alpha \operatorname{vol}(b) = \frac{\alpha}{2^d} \operatorname{vol}(b').$$

Each point in b' can be covered by at most k objects of $\mathcal{U}$, and each large object intersecting b covers a $\alpha/2^d$-fraction of b'. Therefore, there at most $k2^d/\alpha$ objects in $\mathcal{U}$ that intersect g with diameter larger than $2r$. ■

Definition 2.1.7. A metric space $\mathcal{X}$ is a ***doubling space*** if there is a universal constant $c_{\mathrm{dbl}} > 0$ 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***.

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.

Lemma 2.1.8. *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 $\lg = \log_2$ and c_d is the doubling constant of $\mathbb{R}^d$.*

Proof: Cover b by the minimum number of balls of radius $\leq \alpha r$. By the definition of the doubling constant, the number of balls needed is ${c_d}^{\lceil \log_2 1/\alpha \rceil}$. Each of these balls, by definition of density, can intersect at most ρ objects of $\mathcal{U}$ of diameter larger than $2r\alpha$, which implies the claim. ■

The density definition can be made to be somewhat more flexible, as follows.

Lemma 2.1.9. *Let $\beta > 1$ be a parameter, and let $\mathcal{U}$ be a collection of objects in $\mathbb{R}^d$ such that, for any r, any ball with radius r intersects at most ρ objects with diameter $\geq 2r\beta$. Then $\mathcal{U}$ has density ${c_d}^{\lceil \lg \beta \rceil}\rho$.*

Proof: Let b be a ball with radius r. We can cover b with ${c_d}^{\lceil \lg \beta \rceil}$ balls with radius r/β. Each (r/β)-radius ball can intersect at most ρ objects with diameter larger than $2(r/\beta)\beta = 2r$, so b intersects at most ${c_d}^{\lceil \lg \beta \rceil}\rho$ objects with diameter larger than $2r = \operatorname{diam}(b)$. ■

2.1.2. Minors of objects

Definition 2.1.10. A graph G is ***t-shallow*** (or alternatively has *graph radius t*) 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 $h = \operatorname{center}(G)$.

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 2.1). Note that the clusters in Figure 2.1 induce intersection graphs with large graph radius.

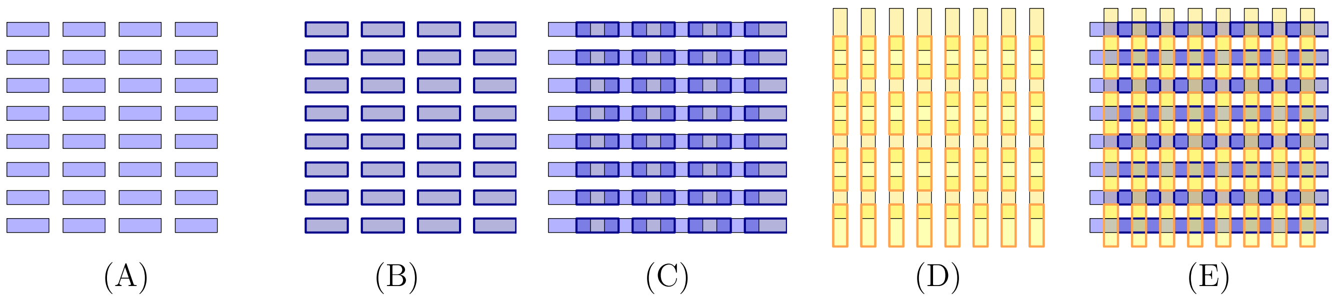

Figure 2.1: (A) and (B) are two low-density collections of n^2 disjoint horizontal slabs, whose intersection graph (C) contains n rows as minors. (D) is the intersection graph of a low-density collection of vertical slabs that contain n columns as minors. In (E), the intersection graph of all the slabs contain the n rows and n columns as minors that form a $K_{n,n}$ bipartite graph, which in turn contains an $n+1$ vertex clique minor.

Definition 2.1.11. For sets of objects $\mathcal{U}$ and $\mathcal{V}$, if $\mathcal{V}$ is a minor of $\mathcal{U}$, and the intersection graph of each cluster of $\mathcal{U}$ 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 2.1.12. *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 its associated cluster $C(g) \subseteq \mathcal{U}$. These sets are disjoint, and let $\mathcal{P} = \{C(g) \mid g \in \mathcal{V}\}$ be the induced partition of $\mathcal{U}$ into clusters (it can also be a partition of a subset of $\mathcal{U}$). Next, consider any ball $b = \mathbb{b}(c, r)$, and suppose that $g \in \mathcal{V}$ intersects b and it has diameter at least $2r$, and let $C(g) \in \mathcal{P}$ be its cluster, and $H = G_{C(g)}$ be its associated intersection graph. By assumption H has (graph) diameter $\leq t$.

Now, 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 $\mathrm{diam}(h') \geq 2r/t$ (if there is no such object then the diameter of $\mathrm{diam}(g) < t(2r/t) \leq 2r$, which is a contradiction).

Consider the shortest path $\pi \equiv h_1, \ldots, h_m$ between $h = h_1$ and $h' = h_m$ in H, where $m \leq t$. Observe that, for $i = 1, \ldots m-1$, $\mathrm{diam}(h_i) < 2r/t$, and thus the distance between b and h' is bounded by $\sum_{i=1}^{m-1} \mathrm{diam}(h_i) \leq (m-1)2r/m < 2r$. We refer to h' as the *representative* of g, denoted by $\mathrm{rep}(g) \in C(g)$.

Now, let $\mathcal{H} = \left\{ \mathrm{rep}(g) \in \mathcal{U} \;\middle|\; g \in \mathcal{V}, \mathrm{diam}(g) \geq 2r, \text{ and } g \cap b \neq \emptyset \right\}$. The representatives in $\mathcal{H}$ are all unique, each is of diameter $\geq 2r/t$, all of them intersect $\mathbb{b}(c, 3r)$, and they all belong to $\mathcal{U}$, a set of density ρ. Lemma 2.1.8 implies that $|\mathcal{H}| \leq \rho {c_d}^{\lceil \lg t \rceil}$. Since $c_d = 2^{O(d)}$, see [Ver05], it follows that $|\mathcal{H}| = t^{O(d)}$, implying the claim. ■

2.2. Graphs with polynomial expansion

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 – based on the edges contraction. 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 2.1.10). Let $\nabla_t(H)$ denote the set of all graphs that are minors[2] of H of depth t.

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

Definition 2.2.1 ([NO08a]). The *greatest reduced average density of rank* r, or just the ***r-shallow density***, of G is the quantity $\mathbb{d}_r(G) = \sup_{H \in \nabla_r(G)} \frac{|E(H)|}{|V(H)|}$.

Definition 2.2.2. The ***expansion*** of a graph class $\mathcal{C}$ is the function $f : \mathbb{N} \to \mathbb{N} \cup \{\infty\}$ defined by $f(r) = \sup_{G \in \mathcal{C}} \mathbb{d}_r(G)$. 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 2.2.3. *If a graph G has bounded expansion, then G has average degree at most μ, where $\mu = \mathbb{d}_1(G)/2 = O(1)$, as the graph G is its own 1-shallow minor, where every vertex is its own cluster. In particular, the vertex v_0 with minimum degree has degree at most μ. Removing v_0 from G leaves a graph G_0 that is also a 1-shallow minor of G, and therefore contains a second vertex v_1 with degree at most μ. Continuing this process, we see that the graph G, by virtue of its bounded expansion, is $O(1)$-degenerate (see Definition 2.1.4).*

As an example of a class of graph 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 2.1.12 together with Observation 2.1.5 implies that low-density graphs have polynomial expansion.

Lemma 2.2.4. *Let $\rho > 0$ be fixed. The graph class $\mathcal{C}^d_\rho$ of ρ-dense graphs in $\mathbb{R}^d$ has polynomial expansion bounded by $f(t) = \rho t^{O(d)}$.*

2.2.1. Separators

Definition 2.2.5. 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 $|Z| = o(|U|)$, and $U \setminus Z$ can be partitioned into two *separate* sets X and Y, with $|X| \leq (2/3)\,|V|$ and $|Y| \leq (2/3)\,|V|$[3].

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 following.

Theorem 2.2.6 ([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 $O\big(mn^{1-\alpha}\log^{1-\alpha} n\big)$ time, a separator of size $O\big(n^{1-\alpha}\log^{1-\alpha} n\big)$, where $\alpha = 1/(2k+2)$.*

For the sake of completeness, a proof is provided in Appendix A.1. This result is well known, and we simply retrace the calculations of [NO08b] for polynomial f instead of subexponential f. Theorem 2.2.6 yields a sublinear separator for low-density graphs of size $O\Big((\rho^2 n \log n)^{1-\frac{1}{O(\log c_{\mathrm{dbl}})}}\Big)$. Geometric arguments give a somewhat stronger separator. For the sake of completeness, we provide a proof in the appendix of the following result, but we emphasize that it is essentially already known [MTTV97, SW98, Cha03].

[3]Here, the choice of 2/3 is arbitrary, and any constant smaller than 1 is sufficient.

Lemma 2.2.7 (Proof in Appendix A.2). *Let $\mathcal{U}$ be a set of n objects in $\mathbb{R}^d$ with density $\rho > 0$ (see Definition 2.1.3$_{p5}$), and let $k \leq n$ be some prespecified number. Then, one can compute, in expected $O(n)$ time, a sphere $\mathbb{S}$ that intersects $O(\rho + \rho^{1/d} k^{1-1/d})$ objects of $\mathcal{U}$. 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.*

2.2.2. Divisions

Consider a set V. A ***cover*** of V is a set $\mathcal{W} = \{C_i \subseteq V \mid i = 1, \ldots, k\}$ 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 2.2.8. 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.

The p 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 $\lambda = \mathrm{poly}(1/\varepsilon)$. Such divisions were first used by Frederickson in planar graphs [Fre87]. For the sake of completeness, we provide a proof of the following well-known result.

Lemma 2.2.9 (Proof in Appendix A.3). *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 $\lambda = O\Big(\big(\varepsilon^{-1} \log^\beta \varepsilon^{-1}\big)^{1/(1-\alpha)}\Big)$.*

Corollary 2.2.10. *(A) Let G be a graph with polynomial expansion of degree k and n vertices, and let $\varepsilon > 0$ be fixed. Then G has $O\big((1/\varepsilon)^{2k+2} \log^{2k+1}(1/\varepsilon)\big)$-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 $\lambda = O\big(\rho/\varepsilon^d\big)$ and total excess at most εn. This division can be computed in $O(n \log(n/\lambda))$ time.

Proof: (A) By Theorem 2.2.6, $\mathcal{C}$ has separator with parameters $\alpha = 1 - 1/(2k+2)$ and $\beta = 1 - 1/(2k+2)$. Plugging this into Lemma 2.2.9 implies λ-divisions where $\lambda = O\Big(\big((1/\varepsilon) \log^\beta (1/\varepsilon)\big)^{1/(1-\alpha)}\Big) = O\big((1/\varepsilon)^{2k+2} \log^{2k+1}(1/\varepsilon)\big)$.

(B) By Lemma 2.2.7, any subgraph of G with m vertices has a separator of size $\leq c\big(\rho + \rho^{1/d} m^{1-1/d}\big)$, for some constant c. Arguing as in Lemma 2.2.9, one can break up G in a a recursive fashion until each portion has size m such that $c\big(\rho + \rho^{1/d} m^{1-1/d}\big) \leq \varepsilon m / c'$, where c' is some absolute constant. As can be easily verified, this holds for $m = \Omega\big(\rho/\varepsilon^d\big)$. Setting $\lambda = m$ implies that the resulting λ-divisions with excess $\leq en$.

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 be done in constant time), using the algorithm of Lemma 2.2.7. ■

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 2.3.1. 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 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).

Lemma 3.1.1. *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 $O\big(n^{O(1)} + 2^\lambda \lambda^{O(1)} n\big)$ time, a set $X \in \Pi$ such that $|X| \geq (1 - \varepsilon) |O|$, where $\lambda = O\Big(\big(\varepsilon^{-1} \log^\beta \varepsilon^{-1}\big)^{1/(1-\alpha)}\Big)$.

Proof: Set $\delta = \varepsilon/2c$. By the algorithmic proof of Lemma 2.2.9, 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 λ. For each cluster, we can find its largest subset with property Π by brute force enumeration in $O\bigl(2^\lambda \lambda^{O(1)}\bigr)$ 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 $O\bigl(n^{O(1)} + 2^\lambda \lambda^{O(1)} n\bigr)$. ∎

Example 3.1.2 (Largest induced planar subgraph). Consider a graph $G = (V, E)$ with n vertices and with polynomial 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 2.2.3, the graph G is t-degenerate, for some $t = O(1)$, since G has a polynomial 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 3.1.1 applies, resulting in an $(1-\varepsilon)$-approximation to the largest induced planar subgraph. The running time of the resulting algorithm is $n^{O(1)} + f(k, \epsilon)n$, for some function f.

Lemma 3.1.3. *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 compute a $(1-\varepsilon)$-approximation to the largest independent set in $\mathcal{U}$. The running time of the algorithm is $O\bigl(n \log n + 2^\lambda \lambda^{O(1)} n\bigr)$, where $\lambda = O\bigl(\rho^{d+1}/\varepsilon^d\bigr)$.*

More generally, one can compute, with the same running time, an $(1-\varepsilon)$-approximate solution for all the problems described in Example 2.3.1.

Proof: Consider the intersection graph $G = G_{\mathcal{U}}$, and observe that by the low-density property, it always have a vertex of degree ρ (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 3.1.1 to G using the improved λ-divisions of Corollary 2.2.10 (B). Here, we need the total excess to be bounded by $(\varepsilon/\rho)n$, which implies that $\lambda = O\bigl(\rho/(\varepsilon/\rho)^d\bigr) = O\bigl(\rho^{d+1}/\varepsilon^d\bigr)$.

For the second part, observe that all the problems mentioned in Example 2.3.1 have solution bigger than the independent set of $\mathcal{U}$, and the same algorithm applies with minor modifications. ∎

Remark 3.1.4. For independent set, one does not need to assume the low density on the input – a more elaborate algorithm works, see Lemma 3.2.2 below.

It is tempting to try and solve problems like dominating set on polynomial-expansion graphs using the algorithm of Lemma 3.1.3. 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 $X \triangle Y = (X \setminus Y) \cup (Y \setminus X)$. Two vertex sets X and Y are ***λ-close*** if $|X \triangle Y| \leq \lambda$; 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 some arbitrary (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. Within a round we can exhaustively try all exchanges in time $n^{O(\lambda)}$, bounding the total running time by $n^{O(\lambda)}$, where n is the size of the ground set of Π.

3.2.3. Analysis of the algorithm

Theorem 3.2.1. *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 for 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 $\mathcal{W} = \{C_1, \ldots, C_m\}$ of K, with boundary vertices B and $\mathrm{excess}(\mathcal{W}) \leq \varepsilon |O \cup L| \leq 2\varepsilon |O|$. For $i = 1, \ldots, m$, let

(i) $O_i = (O \cap C_i) \setminus B, \quad o_i = |O_i|$,
(ii) $L_i = (L \cap C_i), \quad l_i = |L_i|$,
(iii) $B_i = B \cap C_i$, and $b_i = |B_i|$.

Fix i, and consider the set $L' = (L \setminus L_i) \cup O_i$. Since Π is hereditary, $L \setminus L_i \in \Pi$, and since O_i and $L \setminus L_i$ are separated, the set L' is in Π. The set L is λ-locally optimal, and the exchange replacing L_i by O_i can not increase the overall cardinality. We conclude that $l_i \geq o_i$. Summing over all i, we have

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

as desired. ∎

Lemma 3.2.2. *Let $\varepsilon > 0$ and ρ be parameters, and $\mathcal{U}$ 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 for 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 3.2.1, together with the divisions of Corollary 2.2.10 (B), implies the result. ∎

Remark 3.2.3. (A) We emphasize that Lemma 3.2.2 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 2.3.1 have a PTAS using the Lemma 3.2.2 as long as the output has low density.

3.3. Dominating Set

We are interested in approximation algorithm for the following variant of the dominating set problem.

Definition 3.3.1. Let $G = (V, E)$ be an undirected graph, and let D and R be two subsets of V. The set D *dominates* R if every vertex in R either is in D or is adjacent to some vertex in D. 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 task is to compute the smallest subset of D that dominates R.

The algorithm is going to be a local search algorithm, but before analyzing it, we need to develop some tools so we argue about the interaction between the local and optimal solution.

3.3.1. Shallow packings

Definition 3.3.2. Given a graph $G = (V, E)$, a collection of sets $\mathcal{F} = \{C_i \subseteq V \mid i = 1, \ldots, t\}$ 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 2.1.10), 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' are connected in 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 a t-shallow minor.

Lemma 3.3.3. *Let $G = (V, E)$ be an undirected graph, and $\mathcal{F}$ an (ω, t)-cover of G. Then the induced packing graph $H = G[\mathcal{F}]$ has edge density $\frac{|E(H)|}{|V(H)|} \leq 2\omega^2(t+1)^2 \mathrm{d}_t(G) + \omega$, where $\mathrm{d}_t(G)$ is the t-shallow density of G, see Definition 2.2.1.*

Proof: Let the clusters of $\mathcal{F}$ be $\{C_1, \ldots, C_m\}$. 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 $\pi : [\![m]\!] \to [\![m]\!]$ be a random permutation of the cluster indices, chosen uniformly at random, and initialize $\mathcal{F}' = \emptyset$, where $[\![m]\!] = \{1, \ldots, m\}$. For indices $i = 1, \ldots, m$, 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, $C'_{\pi(i)}$ is 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 shrinking randomly the clusters of $\mathcal{F}$.

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 $E_1 = \{C_iC_j \in E(H) \mid c_i \in C_j \text{ or } c_j \in C_i\}$ 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,

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

[4]We allow a set C to appear in $\mathcal{F}$ more than once; that is, $\mathcal{F}$ is a multiset.

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

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 $C'_iC'_j \in E(H)$ 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 $\ell = 2\omega(t+1)$ 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 $C_iC_j \in E_2$, we have $\mathbf{Pr}\big[C'_iC'_j \in E(H)\big] \geq 2/\ell^2$. By linearity of expectation, and since $H' = G[\mathcal{F}']$ is a t-shallow minor of G, we have

$$|E_2| = \sum_{e \in E_2} \frac{\ell^2/2}{\ell^2/2} \leq \frac{\ell^2}{2} \sum_{e \in E_2} \mathbf{Pr}\big[e \in E(H)\big] = \frac{\ell^2}{2} \mathbf{E}\big[|E(H)|\big] \leq \frac{\ell^2}{2} \mathbb{d}_t(G)\,|\mathcal{F}'| \leq \frac{\ell^2}{2} \mathbb{d}_t(G)\,|\mathcal{F}|\,.$$

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

Lemma 3.3.4. *Let G be a graph and $\mathcal{F}$ an (ω, t)-packing of G. Then, for any integer $u > 0$, we have*

$$\mathbb{d}_u\big(G[\mathcal{F}]\big) = 5\omega^2(2u+1)^2(2t+1)^2\mathbb{d}_{2tu+t+u}(G).$$

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

Proof: For $u \geq 1$, a u-shallow minor H of $G[\mathcal{F}]$ is the induced packing graph of a $(\omega, 2tu+u+t)$ cover. By the preceding lemma,

$$\frac{|E(K)|}{|V(K)|} \leq 5\omega^2(4tu+2u+2t+1)^2\mathbb{d}_{2tu+t+u}(G) = 5\omega^2(2t+1)^2(2u+1)^2\mathbb{d}_{2tu+t+u}(G),$$

as desired. ■

3.3.2. Lexical product and shallow density

An interesting consequence of the above is an improvement over known bounds for the shallow density under lexical product (this result is not directly related to the rest of the paper). Given two graphs G and H, the ***lexical 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 3.3.5. *For any graph G, clique K_ω, and $t \in \mathbb{N}$, we have $\mathbb{d}_t(G \bullet K_\omega) \leq 5\omega^2(t+1)^2\mathbb{d}_t(G)$. 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)-cover formed by its clusters. Thus, the claimed inequality follows Lemma 3.3.3. ■

Corollary 3.3.5 is an exponential improvement over the best previously known bounds, on the order of $\mathbb{d}_t(G \bullet K_\omega) \leq \big[O\big(\omega t \mathbb{d}_t(G)\big)\big]^{O(t)}$, 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 3.3.6. For a set of objects $\mathcal{U}$, a collection of subsets $\mathcal{F} = \{C_i \subseteq \mathcal{U} \mid i = 1, \ldots, t\}$ forms a ***(ω, t)-shallow packing*** of G if, for all i, the intersection graph G_{C_i} is t-shallow (see Definition 2.1.10), 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 $\left\{ \bigcup_{f \in C_i} f \;\middle|\; C_i \in \mathcal{F} \right\}$ formed by taking the union of each cluster in $\mathcal{F}$.

Lemma 3.3.7. *Let $\mathcal{U}$ be a collection of objects with density ρ in $\mathbb{R}^d$, and let $\mathcal{F}$ be a (ω, t)-shallow packing. Then the induced object set $\mathcal{U}[\mathcal{F}]$ has density $O(\omega \rho t^{O(d)})$.*

Proof: Consider the collection of objects $\mathcal{V} = \bigsqcup_{C_i \in \mathcal{F}} C_i$ 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 $\omega\rho$. The induced object set $\mathcal{U}[\mathcal{F}]$ is a t-shallow minor of $\mathcal{V}$, so by Lemma 2.1.12, $\mathcal{U}[\mathcal{F}]$ 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 by the vertices that dominate it. In this setting, we prefer the following simple and convenient terminology.

Definition 3.3.8. Given a dominating set $D = \{v_1, \ldots, v_m\}$ 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 $C_1, \ldots, C_m \subseteq D \cup R$ that specifies for every element of D, which elements it covers.

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

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

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 ❀$(D, R) = \{C_1, \ldots, C_m\}$ 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 3.3.9. *Let $G = (V, E)$ be a graph with n vertices and with polynomial expansion of order k, sets $R, D \subseteq V$ such that D dominates R, and let $\varepsilon > 0$ be fixed. Then, for $\lambda = O\big(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon)\big)$, 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.*

Proof: The algorithm starts with the whole collection D as the local solution, and perform legal (and beneficial) local exchanges of size λ until no such exchange is available (see Section 3.2.2), where each local exchange decreases the size of the local solution by one.

Let $O \subseteq D$ and $L \subseteq D$ be the optimal and locally minimum sets dominating R, respectively. Let $\mathcal{O} =$ ❀(O, R) and $\mathcal{L} =$ ❀(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 3.3.4 implies that H has polynomial expansion of order $k + 2$. By Corollary 2.2.10 (A), H has $\lambda = O\big((1/\varepsilon)^{2k+6} \log^{2k+5}(1/\varepsilon)\big)$-division $\mathcal{W} = \{C_1, \ldots, C_m\}$ with a set of boundary vertices B, and total excess $(\varepsilon/4)\,|\mathcal{F}| \le (\varepsilon/4)\big(|\mathcal{O}| + |\mathcal{L}|\big) \le (\varepsilon/2)\,|L|$. For $i = 1, \ldots, m$, 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 $o \in O_i$. As such, the alternative solution L' dominates v, and overall, L' dominates R.

Since L is λ-locally minimal, and the exchange size is $|L \triangle L'| = |L_i \cup O_i| \leq |C_i| \leq \lambda$, implying that L' is at least as large as L. And thus, we have $|L'| = |(L \setminus L_i) \cup O_i| = |L| - |L_i| + |O_i| \geq |L|$. Namely, $|L_i| \leq |O_i|$. Summed over all the clusters W_i, we conclude

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

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

3.3.5. Extensions – multi-cover and reach

One can naturally extend dominating set, as follows:

(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 $\widehat{\delta} = \max_{v \in R} \delta(v)$ 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 $\widehat{\tau} = \max_{v \in R} \tau(v)$.

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 3.3.10. *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 $\widehat{\tau}$ and demand $\widehat{\delta}$ of the given instances are bounded by a constant.*

Then, for $\lambda = O\bigl(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon)\bigr)$, 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 $\mathrm{NN}_k(z, X)$ 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 $\mathrm{NN}_k(z, X)$ 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 $\mathrm{NN}_{\delta(u)}(u, X)$. 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 $\widehat{\tau}$-shallow in G. Let $\mathcal{F} = ❀(X, R) = \{C_v \mid v \in X\}$ be the resulting *flower decomposition* of X.

We claim that a vertex z of G is covered at most $\widehat{\delta}$ times by the flowers of $\mathcal{F}$. More precisely, we prove that z is covered by a flower C_v only if $v \in \mathrm{NN}_{\widehat{\delta}}(z, X)$. For the sake of contradiction, suppose $z \in C_v$ and that $v \notin \mathrm{NN}_{\widehat{\delta}}(z, X)$. Then z is not assigned to v, so there must be a vertex u assigned to z and an associated shortest path $\pi_{uv} = \pi_{uz}|\pi_{zv}$ from u to v through z, where π_{uz} is the subpath from u to z and π_{zu} is the subpath from z to v. Since $v \in \mathrm{NN}_{\widehat{\delta}}(u, X) \setminus \mathrm{NN}_{\widehat{\delta}}(z, X)$, and both sets $\mathrm{NN}_{\widehat{\delta}}(u, X)$ and $\mathrm{NN}_{\widehat{\delta}}(z, X)$ have the same cardinality, there exists another vertex $v' \in \mathrm{NN}_{\widehat{\delta}}(z, X) \setminus \mathrm{NN}_{\widehat{\delta}}(u, X)$. Let $\sigma_{zv'}$ be the shortest-path from z to v'. By construction of $\mathrm{NN}_{\widehat{\delta}}(z, X)$, either $\|\sigma_{zv'}\| < \|\pi_{zv}\|$, or $\|\sigma_{zv'}\| = \|\pi_{zv}\|$ and $v' \prec v$. This implies that either $\|\pi_{uz}|\sigma_{zv'}\| < \|\pi_{uz}|\pi_{zv}\|$, or $v' \prec v$ and $\|\pi_{uz}|\sigma_{zv'}\| = \|\pi_{uz}|\pi_{zv}\|$. In any case, if ties are broken by $\prec$, then v' is closer to u than v is, a contradiction to the premise that $v \in \mathrm{NN}_{\widehat{\delta}}(u, X)$ and $v' \notin \mathrm{NN}_{\widehat{\delta}}(u, X)$. Thus, if z is in a flower C_v, then $v \in \mathrm{NN}_{\widehat{\delta}}(z, X)$.

Now, consider the local solution L and the optimal solution O. Let $\mathcal{O} = ❀(O, R)$ and $\mathcal{L} = ❀(L, R)$ be the flower decompositions of the local and optimal solutions, respectively. Each flower decomposition includes an element at most $\widehat{\delta}$ times, so the combined collection $\mathcal{F} = \mathcal{O} \cup \mathcal{L}$ is a $\left(2\widehat{\delta}, \widehat{\tau}\right)$-shallow packing. By Lemma 3.3.4, the induced packing graph $H = G[\mathcal{F}]$ has polynomial expansion of order $k + 2$, and we can largely follow the argument used in the proof of Theorem 3.3.9. We provide the details for the sake of completeness.

Let $\lambda = O\left(\varepsilon^{-2k+6} \log^{2k+5}(1/\varepsilon)\right)$. There is a λ-division of H into clusters $\mathcal{C}_1, \ldots, \mathcal{C}_m \subseteq \mathcal{F}$, with $\mathcal{B} \subseteq \mathcal{F}$ boundary vertices and total excess $|\mathcal{B}| \leq (\varepsilon/4)\,|\mathcal{F}|$. For $i = 1, \ldots, m$, 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 $v \in V$ such that there is a flower in $\mathcal{L} \setminus \mathcal{L}'$ that covers it (i.e., the vertex "lost" coverage in this potential exchange). This implies that v must be 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 3.3.9. ■

3.3.6. Extension: Connected dominating set

The algorithm of Lemma 3.3.10 can be extended to handle the additional constraint that the computed dominating set is simultaneously connected. In this setting, the local search algorithm only considers beneficial exchanges that result in a connected dominating set.

Lemma 3.3.11. *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 $\widehat{\delta} = \max_{v \in V} \delta(v)$ be bounded by a constant. Then, for $\lambda = O\left(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon)\right)$, 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 use the notations and argument used in Lemma 3.3.10. Here, after a local exchange, the resulting set $L' = (L \setminus L_i) \cup O_i$ might not necessarily connected (although it is still a dominating set).

Let $B_i = \mathrm{heads}(\mathcal{B}_i)$ be the heads of boundary vertices of the ith cluster, see Definition 3.3.8. 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 $\pi_{xy} = \pi_{xv} | \pi_{vy}$, 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 in 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 lies 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 so, 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 $|L'' \triangle L| \leq \lambda'$, 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 3.3.11 extends to constantly bounded reach with an added assumption.

Lemma 3.3.12. *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) $\widehat{\delta} = \max_v \delta(v) = O(1)$ and $\widehat{\tau} = \max_v \tau(v) = O(1)$,

(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 $\lambda = O\big(\varepsilon^{-(2k+6)} \log^{2k+5}(1/\varepsilon)\big)$, 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 3.3.11 goes through, except now the shortest paths between distinct components can be shown to have length at most $2(\widehat{\tau} + 1)$ vertices. Condition (iv) is necessary to keep these paths lying in D. The search size is increased by a factor of $2\widehat{\tau}$ instead of 2, which is only a constant factor difference. ■

3.3.7. Discussion

Observation 3.3.13 (PTAS for vertex cover for graphs with polynomial expansion). *The algorithm of Theorem 3.4.1 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 2.2.2, 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 3.4.1.*

Graphs with subexponential expansion. While we are primarily concerned with graph classes with strongly sublinear separators, the crucial construction of divisions, Lemma 2.2.9, still holds for graph classes with hereditary separators of size $O\left(n/\log^{O(1)} n\right)$. Rather than a $\mathrm{poly}(1/\epsilon)$-division with excess ϵn, we get a $f(1/\epsilon)$-division with excess ϵn for some function f. To this end, one can verify (by the same proof as Theorem 2.2.6) that for a small constant c, if a graph class $\mathcal{C}$ has expansion $\varphi(t) = O\big(\exp(c' \cdot t^{c''})\big)$, for c' and c'' sufficiently small constants, then $\mathcal{C}$ has separators of the desired size $O\left(n/\log^{O(1)} n\right)$. Thus, we obtain PTAS's (with much worse dependence on ϵ) for any graph class $\mathcal{C}$ with subexponential expansion $\mathbb{d}_t(\mathcal{C}) = O\big(\exp(c' \cdot t^{c''})\big)$, where c' and c'' are some constants.

3.4. Geometric applications

The above implies PTAS's for dominating set type problems on low-density graphs. Let $\mathcal{U}$ be a 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 3.3.9 applies because G is low density and as such 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 3.3.9 by applying the stronger separator theorem for low-density graphs.

Theorem 3.4.1. *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 $\varepsilon > 0$ be a parameter. Then, for $\lambda = O\big(\rho/\varepsilon^d\big)$, the local search algorithm, with exchanges of size λ implies the following:*

(A) An approximation algorithm that, in $O\big(mn^{O(\lambda)}\big)$ time, computes a set $Q \subseteq P$ that is an $(1+\varepsilon)$-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 $(1+\varepsilon)$-approximation for the smallest cardinality set that covers P.

Proof: Since points have zero diameter, the union $\mathcal{U} \cup P$ also has density ρ. 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 3.3.9 (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 3.3.9 argues about the induced packing graph of G for some $(2,1)$-shallow packing $\mathcal{G}$. By Lemma 3.3.7, the graph $H = G[\mathcal{G}]$ has density $O(\rho 2^d) = O(\rho)$. Thus, by Corollary 2.2.10 (B), H has a λ-division with excess $(\varepsilon/4)\,|V(H)|$, where $\lambda = O\big(\rho/\varepsilon^d\big)$. The algorithm of Theorem 3.3.9 modified to use these improved divisions implies the result. ∎

Remark 3.4.2. To our knowledge, the algorithms of Theorem 3.4.1 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.

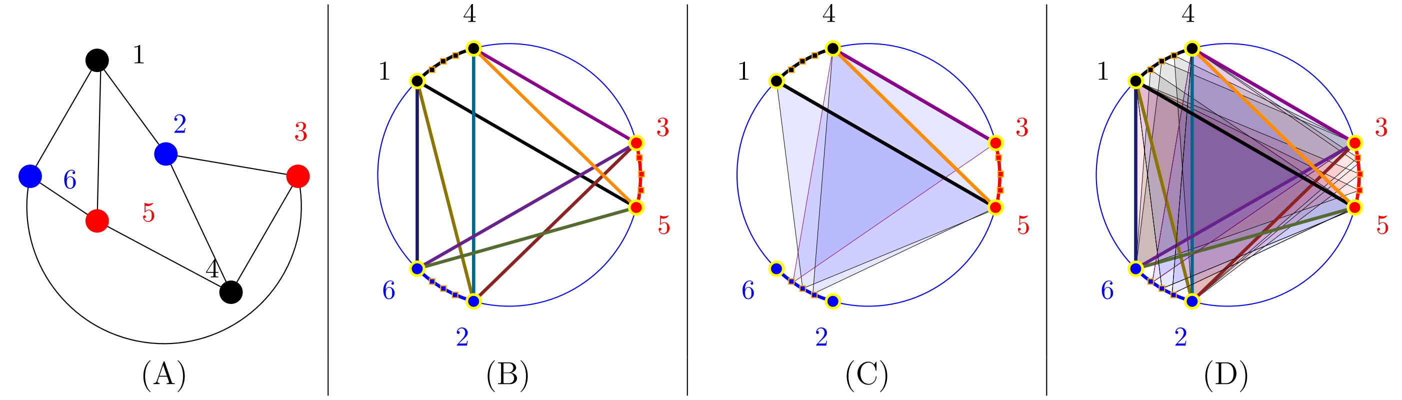

Figure 4.1: Illustration of the proof of Lemma 4.2.1: (A) A 3-regular graph with its 3 coloring. (B) Placing the vertices on a circle. (C) Three edges and their associated triangles. (D) All the triangles.

4. Hardness of approximation

Some of the results of this section appeared in an unpublished manuscript [Har09]. Chan and Grant [CG11] also prove some related hardness results, which were (to some extent) a followup work to the aforementioned manuscript [Har09].

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 kSAT 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 X 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. By the PCP Theorem, this would imply an exact algorithm for 3SAT. Thus, under ETH, showing that a problem is APX-Hard implies that it does not even have a QPTAS, where QPTAS is an $(1+\varepsilon)$-approximation algorithm with running time $n^{O(\mathrm{poly}(\log n, 1/\varepsilon))}$.

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.

Lemma 4.2.1. *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 [CR14][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 $i = 1, \ldots, m$, let $u_i v_i$ be the ith 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 $T_i = \triangle u_i v_i q_i$. 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 $\mathcal{T} = \{T_1, \ldots, T_m\}$ 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
(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.

Lemma 4.3.1. *There is no PTAS for the friendly geometric set cover problem, unless* P = NP*.*

Proof: Let U be a set of n elements, and $\mathcal{F}$ a set of subsets of U each containing at most k elements of U. In the ***minimum k-set cover*** problem, we want to find the smallest subcollection $\mathcal{G} \subseteq \mathcal{F}$ that covers U. The problem is MaxSNP-Hard for $k \geq 3$, meaning there is no PTAS unless P = NP [ACG+99].

We will reduce an instance $(U, \mathcal{F})$ of the minimum k-set cover problem (for $k = 3$) into an instance of the friendly geometric set cover problem.

Let $U = \{u_1, \ldots, u_n\}$ be a set of n elements, and $\mathcal{F} = \{S_1, \ldots, S_m\}$ a collection of m subsets of U. We place n points equally spaced on the unit radius circle centered at the origin, and let $P = \{p_1, \ldots, p_n\}$ 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

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

where $\mathcal{CH}$ is the convex hull, $f(S_i) = \cup_{x \in S_i} \{f(x)\}$, and $\mathsf{disk}(r)$ denotes the disk of radius r centered at the origin. Visually, g_i is a disk with three (since $k = 3$) teeth coming out of it, see Figure 4.2. Note that the boundary of two such shapes intersects in at most 6 points.

It is now easy to verify that the resulting instance of geometric set cover $(P, \{g_1, \ldots, g_m\})$ is friendly, and clearly any cover of P by these shapes can be interpreted as a cover of U by the corresponding sets

[5] Brook'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 $\Delta + 1$.

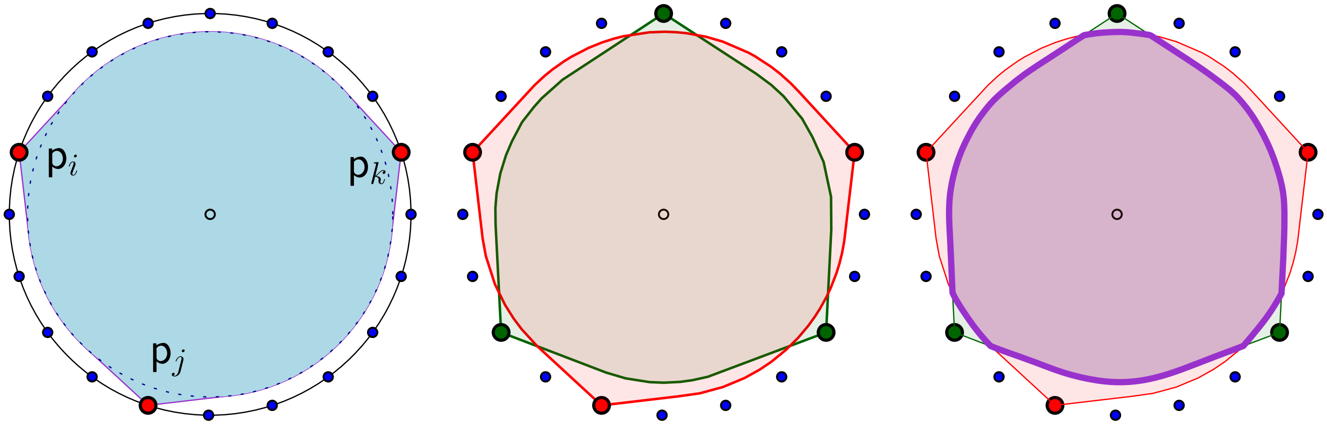

(i) (ii) (iii)

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.

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 $\mathrm{P} = \mathrm{NP}$. ■

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 4.4.1. *There is no PTAS for the fat-triangle set cover problem, unless $\mathrm{P} = \mathrm{NP}$. 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 a graph G with maximum degree three, and observe that a Vertex Cover problem in such a graph can be reduced to Set Cover where every set is of size at most 3. Indeed, the ground set U is the edges of G, and every vertex $v \in V(G)$ gives a rise to the set $S_v = \{uv \in E(G) \mid u \in V(G)\}$, which is of size at most 3. Clearly, any cover C of size t for the set system $\mathcal{X} = \Big(U, \left\{ S_v \,\middle|\, v \in V(G) \right\}\Big)$, has a corresponding vertex cover of G of the same size. Thus, Set Cover with every set of size (at most) three is APX-Hard (this is of course well known). Note that in this set cover instance, every element participates in exactly two sets (i.e., the two vertices adjacent to the original edge).

The graph G has maximum degree three, and by Vizing's theorem [BM76], it is 4 edge-colorable[6]. With regards to the set problem, the ground set of the set system $\mathcal{X}$ can be colored by 4 colors such that no set in this set system has a color appearing more than once.

We are given an instance of the Vertex Cover problem for a graph with maximum degree 3, and we transform it into a set cover instance as mentioned above, denoted by $\mathcal{X} = (U, \mathcal{F}_{\mathcal{X}})$. Let $n = |U|$, and color U (as described above) by 4 colors such that no set of $\mathcal{X}$ has the same color repeated twice, let $U_1, \ldots, U_4$ be the partition of U by the color of the points.

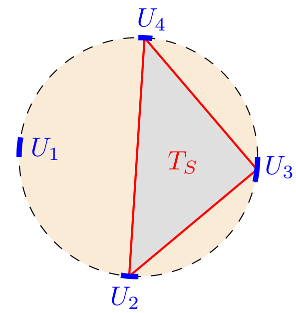

Let $\mathcal{C}$ denote the circle of radius one centered at the origin. We place four relatively short arcs on C, placed on the four intersection points of C with the x and y axes, see figure on the right. Let $I_1, \ldots, I_4$ denote these four circular intervals. We equally space the elements of U_i (as points) on the interval I_i, for $i = 1, \ldots, 4$. 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 $I_1, \ldots, I_4$ 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 4.4.2. For fat triangles of similar size a constant factor approximation algorithm is known [CV07]. Lemma 4.4.1 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 [AdBES14]. It is natural to ask if this can be improved.

4.4.1. Extensions

Lemma 4.4.3. *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 4.4.1, 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 4.4.4. *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.*

[6] Vizing'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.

Proof: Let P be the point set and $\mathcal{F}$ be the set of circles constructed in the proof of Lemma 4.4.3, 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 4.5.1. *There is no PTAS for the maximum independent set of triangles in $\mathbb{R}^3$, unless* $\mathrm{P} = \mathrm{NP}$.

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 $V = \{v_1, \ldots, v_n\}$. We will create a set of triangles, such that their intersection graph is G.

If one spreads n points $p_1, \ldots, p_n$ 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 $i = 1, \ldots, n$.

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 $i = 1, \ldots, n$, 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 $\mathcal{T} = \{f_1, \ldots, f_n\}$. 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 $C_1, \ldots, C_n$ 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 $\mathrm{P} = \mathrm{NP}$. ■

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

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 4.6.1. *Assuming the exponential time hypothesis* (ETH) *(see Section 4.1). Consider a given set of fat triangles $\mathcal{T}$, and a set of points P and density ρ, 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 no $(1+\varepsilon)$-approximation algorithm for the geometric set cover (or hitting set) instance $(P, \mathcal{T})$ with running time $n^{\mathrm{poly}(\log \rho, 1/\varepsilon)}$.

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 $\mathrm{polylog}\, n$ triangles/points and $\mathrm{polylog}\, n$ density, by Lemma 4.4.1. As such, a PTAS in this case, would contradict ETH.

(B) An algorithm with running time $n^{\mathrm{poly}(\log \rho, 1/\varepsilon)}$, would imply via the instance of part (A), that an instance of 3SAT with $\mathrm{polylog}\, n$ number of variables can be solved in $n^{O(\mathrm{poly}(\log \log n))}$ time, contradicting ETH.

The same conclusions holds for geometric hitting set, by using Lemma 4.2.1. ■

5. Conclusions

In this paper, we studied the class of graphs arising out of low density objects in $\mathbb{R}^d$, and showed that the belong to the class of graphs that have polynomial expansion. We also provided PTAS's for independent set and dominating set problems (and some variants) for such graphs.

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. 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 an improved bound for Lemma 2.2.7$_{\text{p9}}$, detailed in the Appendix A.2. We also thank the anonymous referees. 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

[ACG+99] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and approximation*. Springer-Verlag, Berlin, 1999.

[AdBES14] 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.

[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.M. Andreev. On convex polyhedra in lobachevsky spaces. *Sbornik: Mathematics*, 10:413–440, April 1970.

[APS08] P. K. Agarwal, J. Pach, and M. Sharir. State of the union–of geometric objects. In J. E. Goodman, J. Pach, and R. Pollack, editors, *Surveys in Discrete and Computational Geometry Twenty Years Later*, volume 453 of *Contemporary Mathematics*, pages 9–48. Amer. Math. Soc., 2008.

[AW13] A. Adamaszek and A. Wiese. Approximation schemes for maximum weight independent set of rectangles. In *Proc. 54th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), pages 400–409, 2013.

[AW14] A. Adamaszek and A. Wiese. A QPTAS for maximum weight independent set of polygons with polylogarithmic many vertices. In *Proc. 25th ACM-SIAM Sympos. Discrete Algs.* (SODA), pages 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. van Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Santa Clara, CA, USA, 3rd edition, 2008.

[BKSV02] M. de Berg, M. J. Katz, A. F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. *Algorithmica*, 34:81–97, 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. In *Proc. 41st Annu. ACM Sympos. Theory Comput.* (STOC), pages 631–638, 2009.

[CG11] T. M. Chan and E. Grant. Exact algorithms and APX-hardness results for geometric set cover. In *Proc. 23rd Canad. Conf. Comput. Geom.* (CCCG), 2011.

[CG14] S. Cabello and D. Gajser. Simple PTAS's for families of graphs excluding a minor. *CoRR*, abs/1410.5778, 2014.

[CH12] T. M. Chan and S. Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete Comput. Geom.*, 48:373–392, 2012.

[Cha03] T. M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003.

[CR14] D. W. Cranston and L. Rabern. Brooks' theorem and beyond. *ArXiv e-prints*, March 2014.

[CV07] K. L. Clarkson and K. R. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete Comput. Geom.*, 37(1):43–58, 2007.

[dB08] M. de Berg. Improved bounds on the union complexity of fat objects. *Discrete Comput. Geom.*, 40(1):127–140, 2008.

[EK03] J. Erickson and S. Kim. Arbitrarily large neighborly families of congruent symmetric convex 3-polytopes. In A. Bezdek, editor, *Discrete Geometry: In Honor of W. Kuperberg's 60th Birthday*, Lecture Notes Pure Appl. Math., pages 267–278. Marcel-Dekker, 2003.

[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. In *Proc. 20th Annu. ACM Sympos. Theory Comput.* (STOC), pages 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.

[GKS14] M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. In *Proc. 46th Annu. ACM Sympos. Theory Comput.* (STOC), pages 89–98, 2014.

[Gro03] M. Grohe. Local tree-width, excluded minors, and approximation algorithms. *Combinatorica*, 23(4):613–632, 2003.

[Har09] S. Har-Peled. Being fat and friendly is not enough. *CoRR*, abs/0908.2369, 2009.

[Har13] S. Har-Peled. A simple proof of the existence of a planar separator. *CoRR*, abs/1105.0103, 2013.

[Har14] S. Har-Peled. Quasi-polynomial time approximation scheme for sparse subsets of polygons. In *Proc. 30th Annu. Sympos. Comput. Geom.* (SoCG), pages 120–129, 2014.

[Hås96] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proc. 37th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), pages 627–636, 1996.

[HKRS97] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *J. Comput. Sys. Sci.*, 55:3–23, August 1997.

[HQ15] S. Har-Peled and K. Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. In *Proc. 23nd Annu. European Sympos. Algorithms* (ESA), 2015. To appear.

[HR13] S. Har-Peled and B. Raichel. Net and prune: A linear time algorithm for euclidean distance problems. In *Proc. 45th Annu. ACM Sympos. Theory Comput.* (STOC), pages 605–614, New York, NY, USA, 2013. ACM.

[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. In *Complexity of Computer Computations*, pages 85–103, 1972.

[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.

[MRR14b] N. H. Mustafa, R. Raman, and S. Ray. Settling the APX-hardness status for geometric set cover. In *Proc. 55th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), pages 541–550, 2014.

[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*, volume 28 of *Alg. Combin.* Springer, 2012.

[PA95] J. Pach and P. K. Agarwal. *Combinatorial Geometry.* John Wiley & Sons, 1995.

[PRS94] S. Plotkin, S. Rao, and W.D. Smith. Shallow excluded minors and improved graph decompositions. In *Proc. 5th ACM-SIAM Sympos. Discrete Algs.* (SODA), pages 462–470, 1994.

[RS97] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proc. 29th Annu. ACM Sympos. Theory Comput.* (STOC), pages 475–484, 1997.

[Sei91] R. Seidel. Exact upper bounds for the number of faces in d-dimensional Voronoi diagrams. In P. Gritzman and B. Sturmfels, editors, *Applied Geometry and Discrete Mathematics: The Victor Klee Festschrift*, volume 4 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 517–530. Amer. Math. Soc., 1991.

[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, Dept. Comput. Sci., Courant Inst. Math. Sci., New York Univ., New York, NY, 1985.

[Sta92] A. F. van der Stappen. *Motion Planning Amidst Fat Obstacles.* PhD thesis, Utrecht University, Netherlands, 1992.

[SW98] W. D. Smith and N. C. Wormald. Geometric separator theorems and applications. In *Proc. 39th Annu. IEEE Sympos. Found. Comput. Sci.* (FOCS), pages 232–243, 1998.

[Tie05] H. Tietze. Uber das problem der nachbargeibiete 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.

A. Proofs

A.1. Proof of Theorem 2.2.6

Proving Theorem 2.2.6 requires the following result.

Theorem A.1.1 ([PRS94, Theorem 2.3]). *Let G be a graph with m edges and n vertices, and let $\ell, h \in \mathbb{N}$ be two integer parameters. There is an $O(mn/\ell)$ time algorithm that either produces*

(a) the clique K_h as a $\ell \log n$-shallow minor of G, or

(b) a separator of size at most $O(n/\ell + 4\ell h^2 \log n)$.

Restatement of Theorem 2.2.6. *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 $O\big(mn^{1-\alpha}\log^{1-\alpha} n\big)$ time, a separator of size $O\big(n^{1-\alpha}\log^{1-\alpha} n\big)$, where $\alpha = 1/(2k+2)$.*

Proof: Let z be a parameter to be fixed shortly, and let $\ell = z/\log n$ and $cz^k/4 > \mathbb{d}_z(G)$, where c is a sufficiently large constant. Consider a z-shallow minor H of G with $h = cz^k$ vertices, and observe that by definition, we have that $|E(H)| \leq \mathbb{d}_z(G)\,|V(H)| < \dfrac{cz^k}{4}cz^k < \dbinom{h}{2}$. That is, the graph H can not be the clique K_h.

Now, by Theorem A.1.1, G has a separator of size

$$O\Big(n/\ell + \ell h^2 \log n\Big) = O\left(\frac{n\log n}{z} + \frac{z}{\log n}\cdot z^{2k}\cdot \log n\right) = O\left(\frac{n\log n}{z} + z^{2k+1}\right) = O\Big(n^{\frac{2k+1}{2k+2}}\log^{\frac{2k+1}{2k+2}} n\Big)$$

for $z = n^{1/(2k+2)}\log^{1/(2k+2)} n$. The algorithm provided by Theorem A.1.1 runs in time $O\big(\frac{mn}{\ell}\big) = O\big(\frac{mn\log n}{z}\big) = O\big(mn^{\frac{2k+1}{2k+2}}\log^{\frac{2k+1}{2k+2}} n\big)$. ∎

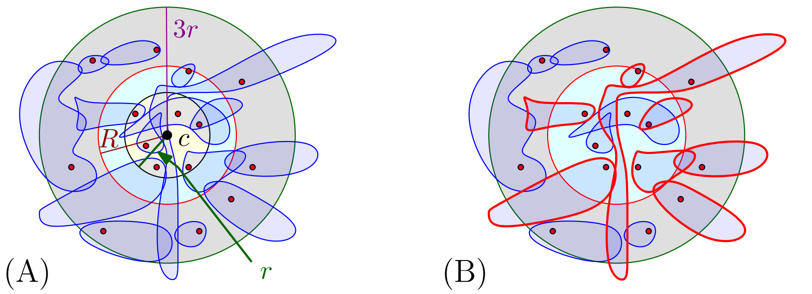

Figure A.1: Illustration of the proof of Lemma 2.2.7.

(A) The ball $\mathbb{b}(c, r)$, and the separating sphere $\mathbb{S}(c, R)$.

(B) All the objects intersecting $\mathbb{S}(c, R)$ are in the separating set.

A.2. Separators for low-density objects

The following result is implied readily by the known results of Miller *et al.* [MTTV97], Smith and Wormald [SW98], and Chan [Cha03]. We provide a proof here for the sake of completeness, and since it is arguably simpler and more elegant.

Restatement of Lemma 2.2.7. *Let $\mathcal{U}$ be a set of n objects in $\mathbb{R}^d$ with density $\rho > 0$ (see Definition 2.1.3$_{\text{p5}}$), and let $k \leq n$ be some prespecified number. Then, one can compute, in expected $O(n)$ time, a sphere $\mathbb{S}$ that intersects $O\big(\rho + \rho^{1/d} k^{1-1/d}\big)$ objects of $\mathcal{U}$. 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.*

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. Next, let $\mathbb{b}(c, r)$ 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 $b = \mathbb{b}(c, R)$ 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. The sphere $\mathbb{S}$ can intersect only $O(\rho/t^{d-1})$ 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 intersect it. Each ball of $\mathcal{B}$ get changed ρ times at most.

Furthermore, any object of $\mathcal{U}$ that intersects $\mathbb{S}$, and has diameter $\leq \ell \leq r$, is fully contained in $b' = \mathbb{b}(c, 3r)$, and b' can be covered by ${c_d}^2$ balls of radius r. As such, b' contains at most $k' = {c_d}^2 k$ points of P, which implies that b' can contain at most $k' = O(k)$ objects of $\mathcal{U}$ inside it, where c_d is the doubling constant of $\mathbb{R}^d$ (see Definition 2.1.7$_{\text{p6}}$). Namely, b' fully contains at most $O(k)$ objects of $\mathcal{U}$ of diameter $\leq \ell$.

Let $\mathcal{U}_{b'}$ be the sets of objects of $\mathcal{U}$, of diameter $\leq \ell$, that are contained in b'. For an 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 as such, it "intersects" $\mathbb{S}$), if $\mathbb{S}$ separates p from q (thus potentially allowing g to be disconnected). Thus, we have that

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

We conclude that the separator size, in expectation, is

$$N = O\Big(\rho + \rho/t^{d-1} + \sum\nolimits_{g \in \mathcal{U}_{b'}} \alpha(g)\Big) = O\Big(\rho + \rho/t^{d-1} + kt\Big)$$

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

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 [HR13]. Using such an approximation slightly deteriorates the constants in the bounds. By Markov's inequality, $\mathbb{S}$ intersects at most $2\alpha'$ objects of $\mathcal{U}$ with probability $\geq 1/2$. If this is not true, we rerun the algorithm. Clearly, in expectation, after a constant number of iterations the algorithm would succeed in finding a sphere that intersects at most $2N$ objects of $\mathcal{U}$. ∎

Remark A.2.1. Mark de Berg (personal communication) pointed out the current simplified proof of Lemma 2.2.7. The authors thank him for pointing out the simpler proof.

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

Lemma A.2.2. *Let $\mathcal{U}$ be a set of n objects in $\mathbb{R}^d$ with density ρ, and weights $\mathrm{w} : \mathcal{U} \to \mathbb{R}$. Let $W = \sum_{f \in \mathcal{U}} \mathrm{w}(f)$ be the total weight of all objects in $\mathcal{U}$. Then one can compute, in expected linear time, a sphere $\mathbb{S}$ that intersects $O\big(\rho + \rho^{1/d} n^{1-1/d}\big)$ 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 2.2.7. We pick a representative point from each object, and assign it the weigh of the object. Next, we compute the smallest ball containing $\geq cW$ of the total weight of the points, and the rest of the proof follows readily, observing that in the worst case, n objects might be involved in the calculations. ∎

A.3. Hereditary separators imply small divisions

Restatement of Lemma 2.2.9. *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 $\lambda = O\Big(\big(\varepsilon^{-1} \log^\beta \varepsilon^{-1}\big)^{1/(1-\alpha)}\Big)$.*

Proof: Our strategy is to break G into smaller pieces. Specifically, at every step the algorithm takes the largest remaining piece $G_{|U}$, compute a balanced separator $Z \subseteq U$ for it, with $L, R \subseteq U$ being the two separated pieces. Specifically, we have

(i) $Z = L \cap R$,
(ii) $L \cup R = U$,
(iii) $|L| \leq (2/3)\,|U|$ and $|R| \leq (2/3)\,|U|$ (see Definition 2.2.5),
(iv) $L \setminus Z$ is separated from $R \setminus Z$ in $G_{|U}$, and
(v) $|Z| \leq f\big(|U|\big)$, where $f(m) \leq c m^\alpha \log^\beta m$, where c is a sufficiently large constant.

Now, the algorithm replaces $G_{|U}$ by the two "broken" pieces $G_{|L}$ and $G_{|R}$. The algorithm continues in this process until all pieces are of size smaller than b (and by construction, of size at least, say, $b/4$), where b is some parameter to be specified shortly. This generates a natural binary separator tree, where the final pieces of the division are the leafs.

Let $N_i = (3/4)^i n$, for $i = 0, \ldots, h = \lceil \log_{4/3} n \rceil$. A piece G_U is at *level* i if $N_{i+1} < |U| \leq N_i$. Consider such a subproblem at node y, which is at level i with ν vertices. The total size of the subproblems of its two children is $\leq \nu + 2f(\nu)$ (here, somewhat confusingly, we count the separator vertices as new, in both subproblems – this makes the following argument somewhat easier). Importantly, each of the subproblems is of size $\leq (2/3)\nu + f(\nu) \leq (3/4)\nu$, implying that both subproblems are in strictly lower level. As such, the fraction of the new vertices created as subproblems move from the ith level to the next is bounded by

$$\nu + 2f(\nu) \leq \nu + 2c\nu^{\alpha}\log^{\beta}\nu = \left(1 + \frac{2c\log^{\beta}\nu}{\nu^{1-\alpha}}\right)\nu \leq \gamma_i \nu,$$

for $\gamma_i = 1 + 2\dfrac{c\log^{\beta} N_{i+1}}{(N_{i+1})^{1-\alpha}}$. In particular, the total number of vertices in the kth level is at most $\Delta_k n$, where

$$\begin{aligned}\Delta_k &= \prod_{j=0}^{k-1}\gamma_j \leq \prod_{j=0}^{k-1}\exp\left(2\frac{c\log^{\beta} N_{j+1}}{(N_{j+1})^{1-\alpha}}\right) = \exp\left(\sum_{j=0}^{k-1}\frac{2c\log^{\beta} N_{j+1}}{(N_{j+1})^{1-\alpha}}\right) \leq \exp\left(\frac{c'\log^{\beta} N_k}{(N_k)^{1-\alpha}}\right)\\ &\leq 1 + \frac{2c'\log^{\beta} N_k}{(N_k)^{1-\alpha}},\end{aligned}$$

since the summation behaves like an increasing geometric series, and c' is a constant that depends on c. The last step follows as $e^x \leq 1 + 2x$, for $0 \leq x \leq 1/2$. In particular, because of the double counting of the separator vertices, the total number of marked vertices in the first k levels is bounded by $n(\Delta_k - 1)$. As such, we need that $\Delta_k - 1 \leq \varepsilon$. This is equivalent to

$$\frac{2c'\log^{\beta} N_k}{(N_k)^{1-\alpha}} \leq \varepsilon \iff \frac{2c'}{\varepsilon} \leq \frac{(N_k)^{1-\alpha}}{\log^{\beta} N_k},$$

which holds if $N_k \geq \left(c''\varepsilon^{-1}\log^{\beta}\varepsilon^{-1}\right)^{1/(1-\alpha)}$, where c'' is a sufficiently large constant. In particular, setting b to (say) twice this threshold implies the claim. ∎