

Chapter 5: Polynomial expansion

Compilation date: January 24, 2020

1 Introduction

In this chapter we are going to study important restrictions of classes of bounded expansion, namely classes of *polynomial expansion*. These are classes \mathcal{C} for which there exists a polynomial $p(x)$ such that $\nabla_r(\mathcal{C}) \leq p(r)$ for all $r \in \mathbb{N}$. The degree of $p(r)$ is the *order* of the class.

Let $\omega_r(\mathcal{C})$ denote the size of the largest clique that can be found as a depth- r minor of a graph from \mathcal{C} . We defined nowhere dense classes by bounding $\omega_r(\mathcal{C})$ for all $r \in \mathbb{N}$, while bounded expansion classes were defined by bounding $\nabla_r(\mathcal{C})$ for all $r \in \mathbb{N}$. Recall that this led to two different notions: there are classes that are nowhere dense but have unbounded expansion. Quite surprisingly, in the setting of polynomial expansion we can equivalently work with $\omega_r(\mathcal{C})$ instead of $\nabla_r(\mathcal{C})$. More precisely, we will prove that there exists a polynomial $p(x)$ such that $\nabla_r(\mathcal{C}) \leq p(r)$ for all $r \in \mathbb{N}$ if and only if there exists a polynomial $q(x)$ such that $\omega_r(\mathcal{C}) \leq q(r)$ for all $r \in \mathbb{N}$, in which case we say that \mathcal{C} has *polynomial ω -expansion*.

Recall again that in the setting of bounded expansion and nowhere denseness we could equivalently work with shallow minors and shallow topological minors, as $\nabla_r(\mathcal{C})$ and $\tilde{\nabla}_r(\mathcal{C})$ are functionally equivalent. However, the upper bound on $\nabla_r(\mathcal{C})$ in terms of $\tilde{\nabla}_r(\mathcal{C})$ had r in the exponent, so this equivalence does not carry over to the setting of polynomial expansion. Indeed, consider the following example. For any fixed integer $d \geq 3$ the class \mathcal{C}_d of all graphs of maximum degree d satisfies $\tilde{\nabla}_r(\mathcal{C}_d) \leq d$ for all $r \in \mathbb{N}$, while it is not hard to prove (say using expanders as examples) that $\nabla_r(\mathcal{C}_d)$ grows exponentially with r .

One of the very useful properties of graph classes with polynomial expansion is the existence of small (sublinear in the graph size) separators. If the class under consideration is closed under taking induced subgraphs we can repeatedly apply the separator theorem to devise efficient divide and conquer algorithms. We have already seen this principle in action when we studied bounded treewidth graphs, which admit balanced separators of constant size. We are going to use the separator properties to devise a polynomial-time approximation scheme (PTAS) for the r -DOMINATING SET problem on such classes.

2 Polynomial expansion and strongly sublinear separators

A *separation* in a graph G is a pair (A, B) of subsets of $V(G)$ such that $A \cup B = V(G)$ and there is no edge between a vertex of $A - B$ and a vertex of $B - A$. The set $A \cap B$ is called the *separator* of the separation and the *order* of the separator is $|A \cap B|$. A separation is *balanced* if $|A - B| \leq \frac{2}{3}|V(G)|$ and $|B - A| \leq \frac{2}{3}|V(G)|$. For a graph class \mathcal{C} and $n \in \mathbb{N}$, let $s_{\mathcal{C}}(n)$ be the smallest integer such that every graph in \mathcal{C} with at most n vertices admits a balanced separation of order at most $s_{\mathcal{C}}(n)$. We say that the class \mathcal{C} has *strongly sublinear separators* if there exists $c \geq 1$ and $0 < \delta \leq 1$ such that

$$s_{\mathcal{C}}(n) \leq c \cdot n^{1-\delta} \quad \text{for every } n \geq 1.$$

Our first result shows that graphs with polynomial ω -expansion admit strongly sublinear separators.

Theorem 1. *Let G be a graph with n vertices and m edges and let ℓ, h be positive integers. There exists an $\mathcal{O}(nm)$ -time algorithm that finds either a complete graph K_h as a depth- $(\ell \log n)$ minor of G or a balanced separation of order at most $\mathcal{O}(n/\ell + \ell h^2 \log n)$ in G .*

Before we prove the theorem, let us see how to derive the desired strongly sublinear bounds for classes of polynomial ω -expansion.

Corollary 2. *Let G be an n -vertex graph from a class with polynomial expansion; more precisely, assume there are constants $d \geq 0$ and $c \geq 1$ such that for all $r \in \mathbb{N}$ we have $\omega_r(G) \leq c(r+1)^d$. Let $\delta = \frac{1}{4d+3}$. Then G admits a balanced separator of order $\mathcal{O}(n^{1-\delta})$. Furthermore, there exists an algorithm that computes such a balanced separator of G in time $\mathcal{O}(mn)$, where $m = |E(G)|$.*

Proof. For $r \in \mathbb{N}$, let $f(r) := c(r+1)^d$. We would like to find functions $\ell: \mathbb{N} \rightarrow \mathbb{N}$ and $h: \mathbb{N} \rightarrow \mathbb{N}$ satisfying the following two conditions:

- (i) $f(\ell(n) \log n) < h(n)$ for sufficiently large n ; and
- (ii) $\mathcal{O}(n/\ell + \ell h^2 \log n) = \mathcal{O}(n^{1-\delta})$.

By the assumptions on the expansion properties of G , condition (i) entails that that K_h is not a depth- $(\ell \log n)$ minor of G . Hence, when applying Theorem 1 to G with $\ell = \ell(n)$ and $h = h(n)$, we can conclude that G has a balanced separator of order $\mathcal{O}(n/\ell + \ell h^2 \log n)$. Condition (ii) then implies that this is a balanced separator of size $\mathcal{O}(n^{1-\delta})$, which can be computed in time $\mathcal{O}(nm)$.

We claim that functions $\ell(n) = \lceil n^\delta \rceil$ and $h(n) = \lceil n^{1/4-\delta/2} \rceil$ satisfy these conditions. We have

$$f(\ell(n) \cdot \log n) = c \left(\lceil n^\delta \rceil \cdot \log n + 1 \right)^d \in \mathcal{O} \left(n^{d\delta} \log^d n \right).$$

Now recall that $\delta = 1/(4d+3)$. We hence have $1/(4d+2) > \delta$, which is equivalent to $1/4 - \delta/2 > d\delta$. This implies condition (i). For condition (ii), observe that

$$\mathcal{O}(n/\ell + \ell h^2 \log n) = \mathcal{O}(n^{1-\delta} + n^{1/2} \log n) = \mathcal{O}(n^{1-\delta}). \quad \square$$

We now come to the proof of Theorem 1. We need one more lemma, which states in a connected graph we may either find a spanning tree of logarithmic depth, or find a roughly balanced separation. For an integer $i \geq 0$ and a set of vertices S , by $N_i[S]$ we denote the set of vertices at distance at most i from S ; in particular $N_i[S] \supseteq S$.

Lemma 3. *Let G be an n -vertex connected graph and let $\ell \geq 1$ be an integer. Then either*

- (1) G admits a spanning tree of depth at most $\lceil 4\ell \log n \rceil + 2$, or
- (2) there exists $S \subseteq V(G)$ satisfying the following conditions:

- (2a) $|N_1[S] - S| \leq |S|/\ell$ and
- (2b) $|N_1[\bar{S}] - \bar{S}| \leq |\bar{S}|/\ell$, where $\bar{S} = V(G) - S$.

Moreover, given G one of the outcomes can be computed in time $\mathcal{O}(m)$, where $m = |E(G)|$.

Proof. Denote $V = V(G)$. Fix an arbitrary vertex $v \in V$ and run breadth-first search from v . If this breadth-first search discovered that all vertices of G are at distance at most $4\lceil\ell \log n\rceil + 2$ from v , then output the corresponding spanning tree of depth at most $4\lceil\ell \log n\rceil + 2$ as output (1).

Suppose otherwise. Let $p = 2\lceil\ell \log n\rceil$. For each $i \geq 0$ let $R_i = N_{2i}[v]$ and let $Q_i = R_{i+1} - R_i$. In other words, we have $R_0 = \{v\}$, $R_{i+1} = N_2[R_i]$, and $Q_i = N_2[R_i] - R_i$. Note that $R_0 \subseteq R_1 \subseteq R_2 \subseteq \dots$. Observe that by the assumption that the depth of the breadth-first search from v is larger than $2p + 2$, we have that sets Q_i for all $0 \leq i \leq p$ are non-empty and $R_{p+1} \neq V$.

Let an index $i \in \{0, \dots, p\}$ be a *forward leap* if the following condition holds:

$$|Q_i| > |R_i|/\ell.$$

Note that this is equivalent to

$$|R_{i+1}| > (1 + 1/\ell)|R_i|.$$

Observe that if among indices i from $\{0, \dots, p\}$ we have k forward leaps, then

$$|R_{p+1}| \geq (1 + 1/\ell)^k \geq 2^{k/\ell},$$

where the last inequality follows from $(1 + 1/\ell)^\ell \geq 2$. Consequently, if we had $k > \ell \log n$, then $|R_{p+1}| > n$, a contradiction. Hence, there are at most $\ell \log n = p/2$ forward leaps.

Similarly, call an index $i \in \{0, \dots, p\}$ a *backward leap* if

$$|Q_i| > |V - R_{i+1}|/\ell,$$

or equivalently

$$|V - R_i| > (1 + 1/\ell)|V - R_{i+1}|$$

Since $R_{p+1} \neq V$, a symmetric reasoning shows that there are at most $\ell \log n = p/2$ backward leaps, for otherwise we would have $|V - R_0| > n$.




Figure 1: Situation for $p = 10$. Each set Q_j consists of two consecutive layers of the breadth-first search from v . Example forward and backward leaps are depicted in blue and red, respectively. Index 5 is the only one that is neither a forward nor a backward leap, hence we set S (light yellow) to be $Q_0 \cup Q_1 \cup \dots \cup Q_4$ plus the first BFS layer from Q_5 .

Now there are $p + 1$ indices in $\{0, \dots, p\}$, out of which there are at most $p/2$ forward leaps and at most $p/2$ backward leaps. Hence, there is an index i that is neither a forward nor a backward leap; observe that such i can be found in linear time. Set $S = N_1[R_i]$; we now verify that S satisfies conditions (2a) and (2b), so it can be returned as outcome (2).

We have $N_1[S] - S = N_2[R_i] - N_1[R_i] \subseteq Q_i$. Since $|Q_i| \leq |R_i|/\ell$ due to i not being a forward leap, we infer that

$$|N_1[S] - S| \leq |Q_i| \leq |R_i|/\ell \leq |S|/\ell,$$

as required. This establishes (2a). Symmetrically, denoting $\bar{S} = V - S$, we have $N_1[\bar{S}] - \bar{S} \subseteq N_1[R_i] - R_i \subseteq Q_i$. Since $|Q_i| \leq |V - R_{i+1}|/\ell$ due to i not being a backward leap, we infer that

$$|N_1[\bar{S}] - \bar{S}| \leq |Q_i| \leq |V - R_{i+1}|/\ell \leq |\bar{S}|/\ell.$$

This establishes (2b). □

We are now ready to prove Theorem 1.

Proof of Theorem 1. Denote $V = V(G)$ and $p = \lceil 4\ell \log n \rceil + 2$. Our algorithm will proceed in rounds maintaining the following objects:

- A depth- p minor model ϕ of a complete graph K_k for some $k < h$. We maintain a property that each branch set of ϕ consists of at most hp vertices. Denote by M the union of all branch sets of ϕ ; note that thus $|M| < h^2p = \mathcal{O}(h^2\ell \log n)$.
- A set R disjoint from M and containing at most $2n/3$ vertices.
- A set X disjoint from M and R , such that $|X| \leq |R|/\ell$ and all neighbors of R are contained in $M \cup X$.

We start with $\phi = R = X = \emptyset$. Let us describe each round of the algorithm.

First, we check whether $|R| \geq n/3$. If this holds, then set $A = R \cup X \cup M$ and $B = V - R$. Observe that we have $A - B = R$, $B - A = V - (R \cup X \cup M)$, $A \cap B = X \cup M$, and (A, B) is a separation in G . Since $n/3 \leq |R| \leq 2n/3$, we also have $|V - (R \cup X \cup M)| \leq 2n/3$, so (A, B) is in fact a balanced separation. Finally, we have

$$|X \cup M| < |X| + h^2p \leq |R|/\ell + h^2p \leq \mathcal{O}(n/\ell + h^2\ell \log n),$$

so the order of this separation is also as requested. Hence (A, B) can be returned as the output of the algorithm.

Assume then that $|R| < n/3$. Next, inspect connected components of $G - (M \cup R \cup X)$ and suppose for a moment that each of them has less than $n/3$ vertices. Iteratively add vertex sets of these connected components to R until its size reaches at least $n/3$. Note that it will not exceed $2n/3$, as each single connected component brings less than $n/3$ vertices to R . Thus we arrive at a situation when R has at least $n/3$ and not more than $2n/3$ vertices, which we already resolved in the previous paragraph.

We are left with the case when there exists a connected component C of $G - (M \cup R \cup X)$ that contains at least $n/3$ vertices. Perform the following operations:

- Modify ϕ by dropping all branch sets that do not have a neighbor in C ; thus M gets modified accordingly as well.
- Modify R by adding to it all vertices reachable from R in $G - (M \cup X)$ (where M is already modified as above). Thus we still have that all neighbors of R are contained in $M \cup X$. Note that no vertex of C got included into R .




Figure 2: Example situation in the proof of Theorem 1. The bottom two panels depict possible outcomes of the round, corresponding to Case 1 and Case 2 below, respectively.

Observe that after the modifications above we still have that $|R| \leq 2n/3$, since R is disjoint with C , which has size at least $n/3$. Again, if after modifications it turned out that $|R| \geq n/3$, then we may find a balanced separator of suitable order as before, so suppose that still it holds that $|R| < n/3$.

Apply Lemma 3 to the component C . We consider two cases depending on the outcome.

Case 1. Lemma 3 produced a spanning tree T of C of depth at most p . For each branch set I of ϕ , select any neighbor u_I of I in C ; such neighbor exists due to the performed modifications. Construct a new branch set J by including the root-to- u_I path in T for each branch set I of ϕ . Since T has depth at most p , it follows that J has radius at most p and in total contains at most hp vertices. Add J to ϕ . If ϕ thus became a depth- p minor model of K_h , then terminate the algorithm and report it. Otherwise all the invariants are maintained and we may proceed to the next round.

Case 2. Lemma 3 produced a suitable vertex subset S of C . Observe that the outcome of Lemma 3 is also valid if we complement it, so by taking $C - S$ instead of S if necessary, we may assume that $|S| \leq |C|/2$. Now add S to R and add all neighbors of S in C into X . By Lemma 3, the number of those neighbors is not larger than $|S|/\ell$, hence the invariant $|X| \leq |R|/\ell$ is maintained in this way. Further, since R before the inclusion of S was disjoint from C and smaller than $n/3$, the size of R after inclusion of S is bounded by

$$|S| + \min(n/3, n - |C|) \leq |C|/2 + \min(n/3, n - |C|) = \min(n/3 + |C|/2, n - |C|/2) \leq 2n/3.$$

Finally, it still holds that all neighbors of R are contained in $M \cup X$ and the minor model ϕ did

not change. We conclude that all invariants are still maintained and hence we may proceed to the next round.

This concludes the description of the algorithm. For the running time analysis, observe that each round can be implemented in time $\mathcal{O}(m)$ and there are at most n rounds in total. \square

We only state the (partial) converse of Theorem 1. Note that we must make the additional assumption that the class under consideration is closed under taking subgraphs.

Theorem 4. *Let \mathcal{C} be a class of graphs which is closed under taking subgraphs and let $c > 0$ and $0 < \delta \leq 1$. If every n -vertex graph in \mathcal{C} has a balanced separator of order at most $cn^{1-\delta}$, then $\nabla_r(\mathcal{C}) \leq c_1 r^{c_2/\delta}$ for appropriately chosen constants c_1, c_2 .*

Corollary 5. *The following conditions are equivalent for a subgraph-closed class of graphs \mathcal{C} :*

- (a) \mathcal{C} has polynomial expansion;
- (b) \mathcal{C} has polynomial ω -expansion;
- (c) \mathcal{C} admits strongly sublinear separators.

3 Approximation algorithms for graphs of polynomial expansion

We now study approximation algorithms for the r -DOMINATING SET problem on graphs of polynomial expansion. It turns out that the existence of strongly sublinear separators gives rise to a robust partition of the vertex set into regions of roughly prescribed size so that the regions communicate with each other only via interfaces sublinear in the sizes. Such a structure is called a λ -division and it is one of key tools in the design of approximation algorithms in well-separable graphs.

λ -divisions. Let us introduce the relevant definitions. A *cover* of a set V is a family \mathcal{X} of subsets $X \subseteq V$, called *regions*, such that $V = \bigcup \mathcal{X}$. A cover of a graph G is a cover of its vertex set with the following property: for every edge uv of G , there is an element of the cover that contains both u and v . The *degree* of a vertex $v \in V$ in \mathcal{X} is the number of regions of \mathcal{X} that contain v . The *total degree* of the cover \mathcal{X} is the sum of degrees over all $v \in V$. The *maximum degree* of the cover is the maximum degree over all $v \in V$. A vertex $v \in V$ is an *interior vertex* of \mathcal{X} if it appears in exactly one region and a *boundary vertex* otherwise. A cover is called a λ -division if every region has size at most λ .

Before we proceed, we note a basic fact about covers: if the total degree is only slightly larger than the size of the ground set, then the total sum of sizes of boundaries of regions is small.

Lemma 6. *Let \mathcal{X} be a cover of a set V such that the total degree of \mathcal{X} is equal to $|V| + p$. For $X \in \mathcal{X}$, let $B(X)$ be set of boundary vertices of X . Then*

$$\sum_{X \in \mathcal{X}} |B(X)| \leq 2p.$$

Proof. If some $u \in V$ has degree k in \mathcal{X} , then it contributes to $\sum_{X \in \mathcal{X}} |B(X)|$ exactly k times if $k > 1$, and 0 times if $k = 1$; this is upper bounded by $2k - 2$ in both cases. Consequently,

$$\sum_{X \in \mathcal{X}} |B(X)| \leq \sum_{u \in V} (2 \cdot \text{degree}_{\mathcal{X}}(u) - 2) \leq 2|V| + 2p - 2|V| = 2p. \quad \square$$

For a cover \mathcal{X} , the difference between the total degree and the cardinality of the ground set is called the *excess* of the cover. Using this term, Lemma 6 says that the sum of sizes of boundaries of regions is bounded in terms of the excess.

We now move to the main point of this section: the existence of strongly sublinear balanced separators implies that by iteratively breaking the graph we can construct λ -divisions with small total degree for λ bounded by a constant.

Lemma 7. *Let G be an n -vertex graph such that every induced m -vertex subgraph of G has a balanced separator of size $\mathcal{O}(m^{1-\delta})$, for some $0 < \delta \leq 1$. Then, for every $\epsilon > 0$, G admits a λ -division with total degree $(1 + \epsilon)n$, where $\lambda \in \mathcal{O}(\epsilon^{-1/\delta})$. Furthermore, such a λ -division can be computed in polynomial time.*

Proof. We iteratively construct a labeled tree T , called a *decomposition tree*, where each node is labeled with an induced subgraph of G , with G being the label of the root. As long as there is a leaf in T labeled $G[U]$ of size larger than λ , for some constant λ to be specified, compute a balanced separation (A, B) for $G[U]$. Let $Z = A \cap B$ be the separator. We have $A \cup B = U$, $|A - Z| \leq 2/3 \cdot |U|$ and $|B - Z| \leq 2/3 \cdot |U|$, $A - Z$ is separated from $B - Z$ in $G[U]$, and $|Z| \leq f(|U|)$, where $f(m) = cm^{1-\delta}$ for some constant c . Now we attach two children labeled $G[A]$ and $G[B]$ to the processed node. The algorithm terminates when all labels $G[U]$ at leaves of the decomposition tree have size at most λ . We return the leaves of the tree as the regions of the division. It is clear that the algorithm returns a cover of G .

We now show that the returned cover has the claimed degree. First observe that each label $G[U]$ at depth i in the separation tree has size at most $(3/4)^i n =: n_i$, provided we set λ large enough. This follows easily by induction: if a label has size m , then each of its children has size at most $(2/3)m + f(m) \leq (2/3)m + cm^{1-\delta} \leq (3/4)m$, where the last inequality holds $m \geq (12c)^{1/\delta}$; hence we should set $\lambda \geq (12c)^{1/\delta}$ to make this step go through. In particular, the depth of the separation tree is bounded by $h := \lceil \log_{4/3} n \rceil$.

The sum of sizes of the children of a node at level i with m vertices is bounded by

$$m + f(m) = m + cm^{1-\delta} = (1 + cm^{-\delta})m \leq (1 + cn_i^{-\delta})m.$$

Hence, in total we have at most

$$\left(\prod_{i=0}^{h-1} (1 + cn_i^{-\delta}) \right) n$$

vertices at the leaf level of the tree. Now we have

$$\prod_{i=0}^{h-1} (1 + cn_i^{-\delta}) \leq \prod_{i=0}^{h-1} \exp(cn_i^{-\delta}) = \exp\left(\sum_{i=0}^{h-1} cn_i^{-\delta}\right) \leq \exp(c'n_h^{-\delta}) \leq 1 + 2c'n_h^{-\delta}.$$

For the for-last inequality, note that $n_i^{-\delta} = (3/4)^{(h-i)\delta} \cdot n_h^{-\delta}$ and $c \cdot \sum_{i=0}^{h-1} (3/4)^{(h-i)\delta}$ is a sum of a decreasing geometric series that can be bounded by a constant c' . In the last inequality we used the known bound $e^x \leq 1 + 2x$ that holds for $0 \leq x \leq 1/2$. Now to achieve the required total degree we need $2c'n_h^{-\delta} \leq \epsilon$, which is equivalent to $(2c'/\epsilon)^{1/\delta} \leq n_h$. Hence we may set the threshold at which we terminate the recursion as

$$\lambda := \max\left((12c)^{1/\delta}, (2c'/\epsilon)^{1/\delta}\right) \in \mathcal{O}(\epsilon^{-1/\delta}). \quad \square$$

Corollary 8. *Let G be an n -vertex graph with polynomial expansion of degree d and let $\epsilon > 0$. Then G has $\mathcal{O}((1/\epsilon)^{4d+3})$ -divisions with total degree $(1 + \epsilon)n$.*

Congested minors. We recall congested shallow minors from Chapter 1.

Definition 1. We say that a graph H is a *depth- d minor with ply c* of a graph G if there exists a mapping ϕ from vertices of H to connected subgraphs of G such that:

- for every $uv \in E(H)$, the subgraph $\phi(u)$ and $\phi(v)$ *touch*: they either share a vertex or there is an edge in G with one endpoint in $\phi(u)$ and second in $\phi(v)$;
- for every $u \in V(H)$, the subgraph $\phi(u)$ has radius at most d ; and
- for every $w \in V(G)$, there are at most c vertices $u \in V(H)$ satisfying $w \in \phi(u)$.

Obviously, depth- d minors of ply 1 are exactly standard depth- d minors. The following is straightforward.

Lemma 9. *Suppose J is a depth- a minor with ply k of H and H is a depth b minor with ply ℓ of G . Then J is a depth- $(2ab + a + b)$ minor with ply $k\ell$ of G .*

Long time ago we did hard work to prove the following.

Lemma 10. *Let G be a graph and let H be a depth- t ply- c minor of G . Then for each $d \in \mathbb{N}$,*

$$\nabla_d(H) \leq c + 2c^2(2dt + d + t + 1)^2 \cdot \nabla_{2dt+d+t}(G).$$

Local search. Local search algorithms are a natural idea in the field of approximation algorithms: we start with any solution and we exhaustively try to make some local improvements. Once no local improvement can be applied, we output the current solution. Local search is very successful in practice, especially in combination with heuristic methods like simulated annealing that enable a broader exploration of the solution space, but is notoriously difficult to analyze in the theoretical (worst-case) setting. We will now use show local search with appropriately large search radius to gives a *polynomial-time approximation scheme* (PTAS) for r -DOMINATING SET on any class of polynomial expansion. Here, a PTAS for a minimization problem is a family of algorithms $(A_\epsilon)_{\epsilon>0}$ such that each A_ϵ computes a $(1 + \epsilon)$ -approximate solution for the problem and runs in polynomial time; the degree of the polynomial may depend in ϵ .

Two sets X and Y are called λ -close if $|X \Delta Y| \leq \lambda$, i.e., if one can transform X into Y by adding and removing at most λ vertices from X . A solution X for some optimization problem is λ -locally optimal if there is no solution Y that is λ -close to X and improves upon X .

Definition 2. The λ -local search algorithm for an optimization problem starts with an arbitrary solution and by examining all λ -close sets, repeatedly makes λ -close improvements until it terminates at a λ -locally optimal solution.

Each improvement in a maximization (minimization) problem increases (resp. decreases) the cardinality of the set, so there are at most n rounds of improvement, where n is the size of the ground set. Within each round we can exhaustively try all λ -close exchanges in time $n^{\mathcal{O}(\lambda)}$, bounding the total running time by $n^{\mathcal{O}(\lambda)}$.

Theorem 11. *Fix $r \in \mathbb{N}$, $\epsilon > 0$, and a class \mathcal{C} of polynomial expansion of order d . Then there is $\lambda \in \mathcal{O}((1/\epsilon)^{4d+11})$ such that the λ -local search algorithm applied to any graph $G \in \mathcal{C}$ computes an r -dominating set of size at most $1 + \epsilon$ times larger than the smallest size of an r -dominating set.*

Proof. Select an arbitrary order \preceq on vertices of G . Let D be an r -dominating set of G . For every vertex $v \in V(G)$, let $\pi(v) \in D$ be the vertex from D that is the closest to v , and in case there are several of them, pick the \preceq -minimal one. Let $P(v)$ be any shortest path from v to $\pi(v)$ in G ; then $P(v)$ has length at most r . For each $u \in D$, define the *cluster* C_u as follows:

$$C_u := \bigcup_{v \in \pi^{-1}(u)} P(v);$$

that is, every $u \in D$ takes to its cluster all the paths $P(v)$ leading to it. Let $\mathcal{D} = \{C_u : u \in D\}$ be the collection of clusters.

Claim 1. *For any r -dominating set D , clusters from \mathcal{D} are pairwise disjoint.*

Proof. For the sake of contradiction suppose some vertex $w \in V(G)$ belongs to two different clusters, say C_u and $C_{u'}$ where $u \prec u'$. This means that there are vertices v and v' such that $\pi(v) = u$, $\pi(v') = u'$, and w belongs to the intersection of $P(v)$ and $P(v')$. Let s be the distance on $P(v)$ from w to u , similarly define s' for $P(v')$ and u' . Observe that it cannot be that $s < s'$, because then v' would be closer to u than to u' , a contradiction with the choice of u' . Similarly it cannot happen that $s' > s$, yielding $s = s'$. But then either v' is closer to u than to u' , or u, u' are equally distant from v' but $u \prec u'$. Again this is a contradiction with the choice of u' . \square

Each cluster from \mathcal{D} obviously has radius at most r , which we will henceforth call the *packing associated with D* .

We choose $\lambda \in \mathcal{O}((1/\epsilon)^{4d+11})$, to be specified later. Consider performing the λ -local search algorithm and suppose it outputs a locally minimum r -dominating set L . Let D be a smallest r -dominating set in G . Let \mathcal{D} and \mathcal{L} be the packings associated with D and L , respectively. Let $\pi_D: V(G) \rightarrow D$ and $\pi_L: V(G) \rightarrow L$ be also the corresponding mappings. For $u \in D$ let $\mathcal{D}(u) = \pi_D^{-1}(u)$ be the unique cluster of \mathcal{D} that contains u , similarly define $\mathcal{L}(u)$.

Let $\mathcal{F} := \mathcal{D} \cup \mathcal{L}$. Note that \mathcal{F} induces a depth- r ply-2 minor of G that we will henceforth denote by H . By Lemma 10, H has polynomial expansion of order $d + 2$. By Lemma 7, there is $\lambda \in \mathcal{O}((1/\epsilon)^{4d+11})$ for which we can find a λ -division \mathcal{X} of H with total degree at most $(1 + \epsilon/16)|\mathcal{F}|$. Note that $|\mathcal{F}| = |\mathcal{D}| + |\mathcal{L}| \leq 2|\mathcal{L}|$, hence the total degree of \mathcal{X} is at most $|\mathcal{F}| + \epsilon/8 \cdot |\mathcal{L}|$.

Now consider any region $X \in \mathcal{X}$; recall that X consists of clusters from \mathcal{F} . Let us define the following sets:

- $\mathcal{B}(X)$ is the set of boundary vertices of X (which are clusters from \mathcal{F}).
- $\mathcal{D}(X) = \mathcal{D} \cap X$ is the set of those clusters in \mathcal{D} that belong to X .
- $\mathcal{L}(X) = \mathcal{L} \cap X - \mathcal{B}(X)$ is the set of those clusters in \mathcal{L} that are interior vertices of X .

Further, let $D(X)$ be the set of those $u \in D$ for which $\mathcal{D}(u) \in \mathcal{D}(X)$, and similarly $L(X)$ is the set of those $u \in L$ for which $\mathcal{L}(u) \in \mathcal{L}(X)$. The next claim is the key point in the proof.

Claim 2. *For any region $X \in \mathcal{X}$, the set $L' := (L - L(X)) \cup D(X)$ is an r -dominating set in G .*

Proof. For the sake of contradiction suppose there exists a vertex $w \in V(G)$ that is not r -dominated by L' . Let $u = \pi_L(w)$ and $v = \pi_D(w)$. Recall that w is r -dominated by u ; since w is not r -dominated by L' , it follows that $u \in L(X)$. Consequently, the cluster $\mathcal{L}(u)$ is not a boundary vertex of X . On the other hand, we have $w \in \mathcal{D}(v)$, hence clusters $\mathcal{L}(u)$ and $\mathcal{D}(v)$ are adjacent in the graph H . Since $\mathcal{L}(u)$ is not a boundary vertex of X , it follows that $\mathcal{D}(v)$ belongs to X . So $v \in D(X)$ and w is r -dominated by $v \in L'$, a contradiction. \lrcorner

Observe that the exchange set described in Claim 2 is of size at most λ , because all exchanged vertices (or rather corresponding clusters) belong to X . Since \mathcal{L} is locally optimal, we obtain that

$$|\mathcal{L}(X)| \leq |\mathcal{D}(X)| \quad \text{for each region } X \in \mathcal{X}.$$

Finally, using this observation and Lemma 6 we estimate the size of L :

$$\begin{aligned} |L| = |\mathcal{L}| &\leq \sum_{X \in \mathcal{X}} |\mathcal{L}(X)| + |\mathcal{B}(X)| \leq \sum_{X \in \mathcal{X}} |\mathcal{D}(X)| + |\mathcal{B}(X)| \\ &\leq |\mathcal{D}| + 2 \sum_{X \in \mathcal{X}} |\mathcal{B}(X)| = |\mathcal{D}| + \epsilon/2 \cdot |\mathcal{L}| = |\mathcal{D}| + \epsilon/2 \cdot |L|. \end{aligned}$$

Hence $(1 - \epsilon)|L| \leq |\mathcal{D}|$, implying $|L| \leq |\mathcal{D}|/(1 - \epsilon/2) \leq (1 + \epsilon)|\mathcal{D}|$. \square