

Optimality of Geometric Local Search

Bruno Jartoux

Laboratoire d'Informatique Gaspard-Monge, Université Paris-Est, ESIEE Paris
Marne-la-Vallée, France
bruno.jartoux@esiee.fr

 <https://orcid.org/0000-0002-5341-1968>

Nabil H. Mustafa

Laboratoire d'Informatique Gaspard-Monge, Université Paris-Est, ESIEE Paris
Marne-la-Vallée, France
mustafan@esiee.fr

Abstract

Up until a decade ago, the algorithmic status of several basic NP-complete problems in geometric combinatorial optimisation was unresolved. This included the existence of polynomial-time approximation schemes (PTASs) for hitting set, set cover, dominating set, independent set, and other problems for some basic geometric objects. These past nine years have seen the resolution of all these problems—interestingly, with the same algorithm: local search. In fact, it was shown that for many of these problems, local search with radius λ gives a $(1 + O(\lambda^{-\frac{1}{2}}))$ -approximation with running time $n^{O(\lambda)}$. Setting $\lambda = \Theta(\epsilon^{-2})$ yields a PTAS with a running time of $n^{O(\epsilon^{-2})}$.

On the other hand, hardness results suggest that there do not exist PTASs for these problems with running time $\text{poly}(n) \cdot f(\epsilon)$ for any arbitrary f . Thus the main question left open in previous work is in improving the exponent of n to $o(\epsilon^{-2})$.

We show that in fact the approximation guarantee of local search cannot be improved for any of these problems. The key ingredient, of independent interest, is a new lower bound on locally expanding planar graphs, which is then used to show the impossibility results. Our construction extends to other graph families with small separators.

2012 ACM Subject Classification Theory of computation \rightarrow Approximation algorithms analysis, Mathematics of computing \rightarrow Combinatorial optimization, Mathematics of computing \rightarrow Graph theory, Mathematics of computing \rightarrow Approximation algorithms

Keywords and phrases local search, expansion, matchings, Hall's marriage theorem

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.48

Funding This work was supported by ANR grant SAGA (JCJC-14-CE25-0016-01).


Acknowledgements We thank the referees for several helpful comments.

1 Introduction

Within the past decade polynomial-time approximation schemes (PTASs) have been proposed for a number of long-standing open problems in geometric approximation algorithms, including the following NP-hard problems (see [18, 10] for hardness results):

Hitting set for pseudodisks [23]: given a set \mathcal{P} of points and a family \mathcal{D} of pseudodisks¹ in the plane, compute a smallest subset of \mathcal{P} that intersects all pseudodisks in \mathcal{D} .

¹ A **family of pseudodisks** is a set of planar regions whose boundaries are Jordan curves and such that the boundaries of any pair of pseudodisks intersect at most twice.



Independent set of pseudodisks [1, 9]: given a family \mathcal{D} of pseudodisks, compute a maximum size subset of pairwise disjoint pseudodisks in \mathcal{D} .

Dominating set of pseudodisks [14, 15]: given a family \mathcal{D} of pseudodisks, compute a smallest subset of pseudodisks of \mathcal{D} that together intersect all other pseudodisks of \mathcal{D} .

Set cover for disks [8, 5]: given a set \mathcal{P} of points and a family \mathcal{D} of disks in the plane, return a smallest subset of disks in \mathcal{D} that together cover all points of \mathcal{P} .

Unit-capacity point-packing [12]: given a set of points \mathcal{P} and a set of disks \mathcal{D} , compute a largest subset of \mathcal{P} that hits no disk of \mathcal{D} more than once.

1.1 Local search

Surprisingly, the PTAS for all these problems is essentially the same: **local search**. Let X be the set of base elements of the problem, and let the **search radius** $\lambda \geq 3$ be an integer. Then start with any feasible solution $\mathcal{S} \subseteq X$ and increase (in the case of a maximisation problem, e.g., maximum independent set) or decrease (in the case of a minimisation problem, e.g., minimum hitting set) its size by local improvement steps while maintaining feasibility. Here a **local improvement step** is to swap a subset \mathcal{S}' of at most λ elements of the current solution \mathcal{S} with a subset of $X \setminus \mathcal{S}$ of size at least $|\mathcal{S}'| + 1$ (for maximisation problems) or at most $|\mathcal{S}'| - 1$ (for minimisation problems), as long as the new solution is still feasible. The algorithm finishes when no local improvement step is possible. Such a solution is called **λ -locally optimal**.

All these algorithms are analysed in a similar way, as follows. Let \mathcal{S} be a λ -locally optimal solution and \mathcal{O} be an optimal solution². To relate the cardinalities of \mathcal{S} and \mathcal{O} , a bipartite **exchange graph** is built on vertex sets \mathcal{S} and \mathcal{O} with a local vertex expansion property³:

Minimisation: for all $\mathcal{S}' \subseteq \mathcal{S}$ of size at most λ , $|N(\mathcal{S}')| \geq |\mathcal{S}'|$. (1)

Maximisation: for all $\mathcal{O}' \subseteq \mathcal{O}$ of size at most λ , $|N(\mathcal{O}')| \geq |\mathcal{O}'|$. (2)

The construction of exchange graphs is problem-specific and exploits the geometric properties of optimal and local solutions. For example, in the minimum vertex cover problem on a graph G this would simply be the bipartite subgraph of G induced by \mathcal{S} and \mathcal{O} ; condition (1) follows from the local optimality of \mathcal{S} .

The key in the analysis lies in a general theorem on local expansion in sparse graphs. A bipartite graph on vertex sets (B, R) is **λ -expanding**⁴ if for all $B' \subseteq B$ of size at most λ we have $|N(B')| \geq |B'|$. A **(vertex) separator** of a graph on n vertices is a subset of vertices whose removal leaves connected components of cardinality at most $\frac{2}{3}n$. A class of graphs \mathcal{G} has the **separator property** with parameter $s \in [0, 1]$ if there exists a positive constant c such that any graph in \mathcal{G} has a separator of size at most cn^{1-s} , where n is the number of vertices. For example, trees have this property with $s = 1$ as they have constant-sized separators, whereas planar graphs have the separator property with parameter $s = \frac{1}{2}$. In fact, the separator property with $s = \frac{1}{2}$ actually holds for graphs excluding fixed minors and in particular for minor-closed classes other than the class of all graphs, e.g. graphs of bounded genus [2]. A class of graphs closed under taking subgraphs is called **monotone**.

² We can assume that these solutions are disjoint by considering $\mathcal{S} \setminus \mathcal{O}$ and $\mathcal{O} \setminus \mathcal{S}$.

³ For a graph $G = (V, E)$ (which will always be clear from the context) and a set $V' \subseteq V$, $N(V')$ denotes the set of neighbours of the vertices of V' in G .

⁴ Note that the roles of B and R are *not* symmetric.

► **Theorem 1** ([23, 9, 5]). *If a finite and λ -expanding bipartite graph on (B, R) belongs to a monotone family with the separator property with parameter $s \in (0, 1)$ and $\lambda \geq \lambda_s$, then $|B| \leq (1 + c_s \lambda^{-s}) \cdot |R|$, where c_s and λ_s are positive constants that depend only on s .*

In an independent paper, Cabello and Gajser [7] describe a subcase of this theorem for K_h -minor-free graphs, which have separators of size $O(h\sqrt{n})$. Finally, Har-Peled and Quanrud [16, 17] observe that intersection graphs of low-density objects in \mathbb{R}^d have the separator property with $s = \frac{1}{d}$.

To complete the analysis for minimisation problems, apply Theorem 1 with $B = \mathcal{S}$ and $R = \mathcal{O}$, and get $|\mathcal{S}| \leq (1 + c_s \lambda^{-s}) \cdot |\mathcal{O}|$. For maximisation problems, take $B = \mathcal{O}$ and $R = \mathcal{S}$, and get $|\mathcal{O}| \leq (1 + c_s \lambda^{-s}) \cdot |\mathcal{S}|$ or equivalently $|\mathcal{S}| \geq (1 - c'_s \lambda^{-s}) \cdot |\mathcal{O}|$.

1.2 Computational efficiency of geometric local search

Given a parameter $\epsilon > 0$, local search with radius $\lambda = \Theta(\epsilon^{-\frac{1}{s}})$ provides a $(1 + \epsilon)$ -approximate solution to problems whose exchange graphs have the separator property with parameter s . This can be implemented in $n^{O(\lambda)}$ time by considering all possible local improvements, thus yielding a PTAS in time $n^{O(\epsilon^{-1/s})}$, and in particular $n^{O(\epsilon^{-2})}$ for the five problems listed on page 1.

The parameterised versions of these problems are W[1]-hard: even for unit disks, independent set is W[1]-complete [20] and dominating set is W[1]-hard [21], and dominating set of unit disks is easily reduced to our other three problems. Under the common assumption that $\text{FPT} \subsetneq \text{W}[1]$, which follows from the exponential time hypothesis, these problems do not admit PTASs with time complexity $\text{poly}(n) \cdot f(\epsilon)$ for any arbitrary function f . In other words, the dependence of the exponent of n on ϵ is inevitable.

Still, this running time is prohibitively expensive, and there have been two complementary approaches towards further progress: firstly, careful implementations of local search that find local improvements more efficiently than by brute force [6]. The second, more structural approach is to better analyse the quality of solutions resulting from local search algorithms, mainly by studying the properties of exchange graphs [3].

1.3 Contributions: tightness of Theorem 1

The construction given in section 2 shows that Theorem 1 is asymptotically tight whenever $\frac{1}{s}$ is an integer.

► **Theorem 2.** *Given a positive integer d , there are positive constants c_d and λ_d such that, for every integer $\lambda \geq \lambda_d$, there is a family of bipartite graphs $(B_n, R_n; E_n)_{n \in \mathbb{N}}$ that*

- *are λ -expanding,*
- *have the separator property for $s = \frac{1}{d}$, and so do their subgraphs,*
- *satisfy $|B_n|, |R_n| = \Theta(n)$ and $|B_n| \geq (1 + c_d \cdot \lambda^{-\frac{1}{d}})|R_n| - o(|R_n|)$ as $n \rightarrow \infty$.*

Furthermore when $d = 2$ they are Gabriel graphs.

(A graph (V, E) is called **Gabriel** if there exists a mapping f of the vertices of V to points in the plane such that $\{v_i, v_j\} \in E$ if and only if the circumdisk of the segment $f(v_i)f(v_j)$ contains no other point of $f(V)$. Gabriel graphs are subgraphs of Delaunay triangulations and thus planar.)

► **Remark.** Since our construction for $d = 2$ is planar, previous analogues of Theorem 1 restricted to planar graphs are also tight.

1.4 Algorithmic consequences

The analysis of local search in terms of the radius that achieves a $(1 + \epsilon)$ -approximation is tight for the five problems listed earlier (which all had $s = \frac{1}{2}$), as well as for a few other problems with small separators (section 3).

► **Theorem 3** ([23]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the minimum hitting set problem for pseudodisks.*

► **Corollary 4.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there is a set \mathcal{D} of at least n disks and two disjoint sets B and R of at least n points in \mathbb{R}^2 each such that both B and R are hitting sets for \mathcal{D} , $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and B is a λ -locally optimal solution to the hitting set problem for \mathcal{D} with $\mathcal{P} = B \cup R$.*

► **Theorem 5** ([9]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the maximum independent set problem for pseudodisks.*

► **Corollary 6.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two independent sets B and R of at least n disks in \mathbb{R}^2 such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally optimal solution to the independent set problem in $B \cup R$.*

► **Theorem 7** ([5, 8]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the minimum set cover problem for disks.*

► **Corollary 8.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two independent sets B and R of at least n disks in \mathbb{R}^2 and a set \mathcal{P} of $\Theta(|R|)$ points in \mathbb{R}^2 such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally optimal solution to the set cover problem for \mathcal{P} in $B \cup R$.*

► **Theorem 9** ([14, 15]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the minimum dominating set problem for pseudodisks.*

► **Corollary 10.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there is a set \mathcal{D} of disks in \mathbb{R}^2 and two dominating sets B and R of \mathcal{D} of at least n disks each such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and B is a λ -locally optimal solution to the dominating set problem for \mathcal{D} .*

► **Theorem 11** ([12]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the maximum unit-capacity point-packing problem for disks.*

► **Corollary 12.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two sets B and R of at least n points in \mathbb{R}^2 and a set \mathcal{D} of $\Theta(|R|)$ disks in \mathbb{R}^2 such that every disk of \mathcal{D} contains one point from B and one point from R , $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally optimal solution to the unit-capacity point-packing problem for \mathcal{D} in $B \cup R$.*

Borrowing a term from Arya et al. [4] we could say that: ‘The locality gap for independent set of disks, dominating set of disks, etc. is $1 + \Theta(\lambda^{-1/2})$ ’.

2 Proof of Theorem 2

In this section we build a family of graphs that have the properties stated in Theorem 2. Namely, given parameters d , a large enough λ and n , we construct a bipartite graph G with vertex set (B, R) such that:

- (a) $|R| = n + o(n)$ as $n \rightarrow +\infty$,
- (b) G is λ -expanding,
- (c) $|B| \geq (1 + c\lambda^{-\frac{1}{d}}) \cdot |R| - o(|R|)$ as $n \rightarrow +\infty$, where c depends only on d ,
- (d) any subgraph of G on m vertices has a separator of size $O(m^{1-\frac{1}{d}})$, and
- (e) G is a Gabriel graph when $d = 2$.

The vertices of R are called the **red vertices**, and the vertices of B the **blue vertices**. Our construction is geometric, in that vertices correspond to points in \mathbb{R}^d . Thus we use the terminology vertex and point interchangeably. We denote the i -th coordinate of a point $p \in \mathbb{R}^d$ by $x_i(p)$.

Let $L \geq 2$ and t be two positive integers whose values will be fixed later as a function of the parameters d , λ and n . Let Ξ be a $L \times \dots \times L$ regular integer grid in \mathbb{R}^d consisting of the $(L + 1)^d$ points in $\{0, \dots, L\}^d$. It has L^d **cells**, each defined by precisely 2^d vertices of Ξ . In every cell of Ξ , the vertex with the lexicographically minimum coordinates among the 2^d red vertices defining it is called the **anchor vertex** of that cell. Each vertex—apart from those with one of the d coordinate values equal to L —is the anchor vertex of exactly one cell, which is called its **top cell**. The cell with anchor vertex $(0, \dots, 0)$ is called the **lowest cell** of Ξ .


We define a first bipartite graph $G(d, L)$ as follows. The red vertices of $G(d, L)$ consist of the $(L + 1)^d$ points of Ξ . We next place a blue vertex at the centre of each of the L^d cells of Ξ —except for the lowest cell, which contains two blue vertices with coordinates $(\frac{1}{4}, \dots, \frac{1}{4}, \frac{3}{4})$ and $(\frac{3}{4}, \dots, \frac{3}{4}, \frac{1}{4})$. Thus $G(d, L)$ has precisely $L^d + 1$ blue vertices. The edges of $G(d, L)$ consist of 2^d edges from each blue vertex to the 2^d red vertices of its cell. Of the two blue vertices in the lowest cell of Ξ , one is connected to all the red vertices of the cell except for $(0, \dots, 0, 1)$ (the vertex v that has $x_i(v) = 1$ if and only if $i = d$) and the other to all red vertices except for $(1, \dots, 1, 0)$.

Our second and final graph $G(d, L, t) = (B, R; E)$ is defined as a $t \times \dots \times t$ grid composed of t^d **translates** of $G(d, L)$. Each translate of $G(d, L)$ is indexed by a vector $\vec{r} \in \{0, \dots, t - 1\}^d$, where by $G^{\vec{r}}$ we denote the translate of $G(d, L)$ by $L \cdot \vec{r}$. The blue vertices of $G(d, L, t)$ are simply the disjoint union of the blue vertices of each $G^{\vec{r}}$; the red vertices are also the union of the red vertices of each $G^{\vec{r}}$, except that we identify duplicate red vertices shared by the boundary of two adjacent grids. See Figure 1 for an example for the case $d = 2$.


An explicit description of $G(d, L, t) = (B, R; E)$ is:

$$\begin{aligned}
 R &= \{0, \dots, tL\}^d, \\
 B &= \left\{ \left(\frac{1}{2}, \dots, \frac{1}{2} \right) + \vec{x} : \vec{x} \in \{0, \dots, tL - 1\}^d \setminus \{0, L, \dots, (t - 1)L\}^d \right\} \\
 &\quad \cup \left\{ (\mu, \dots, \mu, 1 - \mu) + L \cdot \vec{x} : \vec{x} \in \{0, 1, \dots, t - 1\}^d, \mu \in \left\{ \frac{1}{4}, \frac{3}{4} \right\} \right\}, \\
 E &= \left\{ \{b, r\} : (b, r) \in B \times R, \min_{i \in \{1, \dots, d\}} |x_i(b) - x_i(r)| \leq \frac{1}{2}, \max_{i \in \{1, \dots, d\}} |x_i(b) - x_i(r)| \leq 1 \right\}.
 \end{aligned}$$

The $L^d + 1$ blue vertices of $G^{\vec{r}}$ form the set $B_{\vec{r}}$. For the red vertices, note that the outer red vertices of each copy of $G(d, L)$ may be shared between up to 2^d translates. To avoid



■ **Figure 1** The graph $G(d, L)$ (shown on the left for $d = 2$ and $L = 3$) has L^d grid cells. It is the basic building block of the graph $G(d, L, t)$ (right, with $t = 5$). Square vertices are red, round vertices are blue.



■ **Figure 2** The three-dimensional lowest cell of $G(3, L)$.

this overlap, let $R_{\vec{\tau}}$ consist only of the L^d red vertices $v \in G^{\vec{\tau}}$ such that $x_i(v) < L(\vec{\tau}_i + 1)$ for each i . In two dimensions, this amounts to peeling off the $2L + 1$ red vertices located on the top and right boundaries of $G^{\vec{\tau}}$.

Let R_b be the set of red vertices with at least one coordinate value equal to tL . We have

$$B = \bigcup_{\vec{\tau}} B_{\vec{\tau}} \quad \text{and} \quad R = R_b \cup \bigcup_{\vec{\tau}} R_{\vec{\tau}},$$

where all unions are disjoint. Observe that

$$|B| = t^d(L^d + 1) \quad \text{and} \quad |R| = (tL + 1)^d. \tag{3}$$

Local expansion

To prove that $G(d, L, t)$ is locally expanding we fix a subset B' of B and let $R' = N(B')$ be the set of its (red) neighbours in $G(d, L, t)$. We show that $|R'| \geq |B'|$ whenever B' is smaller than some function of L and d ; later we will set L such that this function turns out to be at least λ .

A grid cell is **non-empty** if it contains a vertex of B' and otherwise **empty**. A vertex of R' that belongs to R_b or whose top cell is empty is called a **boundary vertex**.

We first sketch a proof in two dimensions based on a *charging* argument (a one-to-one mapping from B' to R'): each vertex of B' is charged to a vertex of R' such that each vertex of R' receives at most one charge, implying that $|R'| \geq |B'|$. Charge each blue vertex of B' to the anchor red vertex of its cell. For those $G^{\vec{\tau}}$ containing two blue vertices in the

lowest cell, one of them remains uncharged. On the other hand, each red vertex receives one charge, except the boundary vertices which receive zero charge. Now for each $\vec{\tau}$ for which $G^{\vec{\tau}}$ contains at least two boundary red vertices charge the uncharged blue vertex in $G^{\vec{\tau}}$ (if it exists) to any one of these (at least two) boundary vertices.

There still remains an uncharged blue vertex in those $G^{\vec{\tau}}$ with less than two boundary red vertices. However, for each such $\vec{\tau}$, the number of vertices of B' in $G^{\vec{\tau}}$ must be at least $\frac{L^2}{2}$. Thus overall, there can remain at most $\frac{|B'|}{L^2/2} = \frac{2|B'|}{L^2}$ uncharged blue vertices. On the other hand, we argue that the total number of boundary red vertices is at least $c_2 \cdot \sqrt{|B'|}$, for some constant c_2 . By our charging scheme, at least half of them—or $\frac{c_2}{2} \cdot \sqrt{|B'|}$ —are still uncharged. Thus when $\frac{2|B'|}{L^2} \leq \frac{c_2}{2} \cdot \sqrt{|B'|}$ —or equivalently $|B'| \leq c' \cdot L^4$, the number of uncharged blue vertices will be less than the number of uncharged red vertices, and we are done.

Now we present the complete proof for general d . We need two preliminary statements. Let the indicator variable $d_{\vec{\tau}}$ be 1 if both blue vertices in the lowest cell of $G^{\vec{\tau}}$ belong to B' and 0 otherwise. Also let $\delta_{\vec{\tau}}$ be the number of boundary vertices in $R_{\vec{\tau}}$. The total number of boundary vertices in R' is thus

$$\delta = |R_b \cap R'| + \sum_{\vec{\tau}} \delta_{\vec{\tau}}. \tag{4}$$

► **Lemma 13.** *For each index $\vec{\tau}$, if $d_{\vec{\tau}} = 1$ and $\delta_{\vec{\tau}} < 2$, then $|B' \cap B_{\vec{\tau}}| \geq \frac{L^d}{2}$.*

Proof. For such an index, B' contains both blue vertices from the lowest cell of $G^{\vec{\tau}}$ so R' contains the 2^d red vertices of this cell. If $\delta_{\vec{\tau}} = 0$, that is, $R_{\vec{\tau}}$ contains no boundary vertex, then the blue vertex in each other cell of $G^{\vec{\tau}}$ is present in B' , and so B' includes all of $B_{\vec{\tau}}$, which consists of $L^d + 1$ blue vertices. It remains to consider the case when $R_{\vec{\tau}}$ contains one unique boundary vertex $v_r \in R' \cap R_{\vec{\tau}}$.

Without loss of generality, assume that $\vec{\tau} = (0, \dots, 0)$. As both blue vertices from the lowest cell of $G^{\vec{\tau}}$ belong to B' , the boundary vertex v_r cannot be the lowest vertex of $G^{\vec{\tau}}$, which has coordinates $(0, \dots, 0)$. Thus there must be some $j \in \{1, \dots, d\}$ for which $x_j(v_r) > 0$. Consider the grid slab Ξ' consisting of all cells whose anchor vertex v has $x_j(v) = 0$. Note that Ξ' contains the lowest cell of $G^{\vec{\tau}}$, which has two vertices of B' . Thus no other cell of Ξ' can be empty, as otherwise that would imply the existence of another boundary red vertex anchoring one of the cells of Ξ' . Now take any cell c of Ξ' whose anchor vertex differs in at least one coordinate other than x_j from v_r ; there are $L^{d-1} - 1$ such cells. All the L cells of $G^{\vec{\tau}}$ whose anchor vertex only differs in the j -th coordinate value from the anchor vertex of c must also be non-empty, as otherwise it would imply the existence of a boundary red vertex in one of these L cells.

Thus there are at least $L(L^{d-1} - 1)$ non-empty cells in $G^{\vec{\tau}}$, i.e., $|B' \cap B_{\vec{\tau}}| \geq L^d - L$ which is at least $\frac{L^d}{2}$ since $L \geq 2$. ◀

Let T be the set of indices $\vec{\tau}$ with $d_{\vec{\tau}} = 1$ and $\delta_{\vec{\tau}} < 2$. As a consequence of the previous lemma, for every such $\vec{\tau} \in T$, the translate $G^{\vec{\tau}}$ contains at least $\frac{L^d}{2}$ vertices of B' , and thus $|T| \leq 2|B'|L^{-d}$. Now consider the quantity $d_{\vec{\tau}} - \frac{\delta_{\vec{\tau}}}{2}$. If $\vec{\tau} \in T$, we have $d_{\vec{\tau}} = 1$ and $0 \leq \delta_{\vec{\tau}} < 2$ and so $d_{\vec{\tau}} - \frac{\delta_{\vec{\tau}}}{2}$ is at most 1. Otherwise for any $\vec{\tau} \notin T$, it is 0 or negative. Therefore

$$\sum_{\vec{\tau}} \left(d_{\vec{\tau}} - \frac{\delta_{\vec{\tau}}}{2} \right) \leq |T| \leq \frac{2|B'|}{L^d}. \tag{5}$$

An isoperimetric inequality

Consider the set S of all grid cells containing vertices of B' . As each cell contains at most two blue vertices, $|B'| \leq 2|S|$. In the orthogonal projection along the i -th coordinate, S is sent to a set S_i of $(d-1)$ -dimensional cells. The preimage of each cell of S_i is a column of d -dimensional cells and must contain at least one boundary vertex, so the total number δ of boundary vertices is at least $|S_i|$. The combinatorial Loomis–Whitney inequality [19] relates d - and $(d-1)$ -dimensional volumes:

$$\prod_{i=1}^d |S_i| \geq |S|^{d-1} \geq \left(\frac{|B'|}{2}\right)^{d-1},$$

from which it follows that

$$\delta^d \geq \left(\frac{|B'|}{2}\right)^{d-1}. \quad (6)$$

Now we come to the key claim, which means that the graph $G(d, L, t)$ is $(2^{1-3d}L^{d^2})$ -expanding.

► **Lemma 14.** *If $2^{3d-1}|B'| \leq L^{d^2}$, then $|R'| \geq |B'|$.*

Proof. For every index $\vec{\tau}$, by definition, each vertex in the set $R_{\vec{\tau}} \cap R'$ either has its top cell non-empty or is a boundary vertex of $G^{\vec{\tau}}$. The number of non-empty top cells in $G^{\vec{\tau}}$ is $|B_{\vec{\tau}} \cap B'| - d_{\vec{\tau}}$, while the number of boundary vertices is $\delta_{\vec{\tau}}$. Thus

$$\begin{aligned} |R'| &= |R_b \cap R'| + \sum_{\vec{\tau}} |R_{\vec{\tau}} \cap R'| = |R_b \cap R'| + \sum_{\vec{\tau}} (|B_{\vec{\tau}} \cap B'| - d_{\vec{\tau}} + \delta_{\vec{\tau}}) \\ &= |R_b \cap R'| + |B'| - \sum_{\vec{\tau}} (d_{\vec{\tau}} - \delta_{\vec{\tau}}) \\ &\geq |B'| - \sum_{\vec{\tau}} \left(d_{\vec{\tau}} - \frac{\delta_{\vec{\tau}}}{2}\right) + \frac{1}{2} \left(|R_b \cap R'| + \sum_{\vec{\tau}} \delta_{\vec{\tau}}\right) \\ &= |B'| - \sum_{\vec{\tau}} \left(d_{\vec{\tau}} - \frac{\delta_{\vec{\tau}}}{2}\right) + \frac{\delta}{2}. \end{aligned}$$

Use the lower bounds (5) and (6) for the second and third summands:

$$|R'| \geq |B'| - \frac{2|B'|}{L^d} + \frac{1}{2} \left(\frac{|B'|}{2}\right)^{(d-1)/d},$$

and the result $|R'| \geq |B'|$ follows when


$$\begin{aligned} \frac{2|B'|}{L^d} &\leq \frac{1}{2} \left(\frac{|B'|}{2}\right)^{(d-1)/d} \\ 2^{3-1/d}|B'|^{1/d} &\leq L^d \end{aligned}$$

or equivalently $2^{3d-1}|B'| \leq L^{d^2}$. ◀

instances consist of proper disks rather than just families of pseudodisks, the bound applies also to the restrictions of these problems to disk families.

For $d = 2$ and any given $\lambda \geq \lambda_d$ and n , let $G = (B, R; E)$ be the planar and λ -expanding graph $(B_n, R_n; E_n)$ described in Theorem 2 and built in section 2. Our instances are based on G : its vertex sets are associated with feasible solutions of the problems. It then suffices to check that the solution associated with B_n (for minimisation problems) or R_n (for maximisation problems) is locally optimal.

3.1.1 Hitting set for pseudodisks



■ **Figure 4** A ‘tight’ instance for the hitting set problem.

► **Theorem 3** ([23]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the minimum hitting set problem for pseudodisks.*

► **Corollary 4.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there is a set \mathcal{D} of at least n disks and two disjoint sets B and R of at least n points in \mathbb{R}^2 each such that both B and R are hitting sets for \mathcal{D} , $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and B is a λ -locally optimal solution to the hitting set problem for \mathcal{D} with $\mathcal{P} = B \cup R$.*

Proof. Recall that the circumdisk of each edge of G contains only its two endpoints. The input consists of all such disks, with $\mathcal{P} = B \cup R$, so that the hitting sets are exactly the vertex covers of G . By construction both B and R are feasible solutions.

On this instance, a λ -local improvement for B would remove a set B' of blue vertices with $|B'| \leq \lambda$. To preserve the hitting set property, it would then need to add to the solution the red endpoints of all edges with their blue endpoint in B' , i.e. the set $N(B')$. Because the graph is λ -expanding, there are at least $|B'|$ such red neighbours: B is λ -locally optimal. ◀

3.1.2 Independent set of pseudodisks


► **Theorem 5** ([9]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the maximum independent set problem for pseudodisks.*

► **Corollary 6.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two*

independent sets B and R of at least n disks in \mathbb{R}^2 such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally optimal solution to the independent set problem in $B \cup R$.

Proof. Realise the graph G as an intersection graph of red and blue disks. (Because it is planar, the disks could even be taken interior-disjoint by the Koebe–Andreev–Thurston theorem.) The independent sets of disks correspond to the independent sets of G . Since G is bipartite both the blue and red families of disks form independent sets, and the red solution is $(\lambda - 1)$ -locally optimal—in maximisation terms: a $(\lambda - 1)$ -local improvement for the red solution would remove a set R' of up to $\lambda - 1$ red disks and replace them with a set B' of blue disks such that $N(B') \subseteq R'$ (to preserve independence) and $|B'| > |R'|$. If there exists a subset $B'' \subseteq B'$ of size $|R'| + 1$, which is at most λ , then since G is λ -expanding such a set has $|B''| \leq |N(B'')| \leq |R'|$, a contradiction. Thus R is a $(\lambda - 1)$ -locally optimal solution. ◀

3.1.3 Set cover for disks



■ **Figure 5** ‘Tight’ instances for independent set and set cover with disks.

► **Theorem 7** ([5, 8]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the minimum set cover problem for disks.*

► **Corollary 8.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two independent sets B and R of at least n disks in \mathbb{R}^2 and a set \mathcal{P} of $\Theta(|R|)$ points in \mathbb{R}^2 such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally optimal solution to the set cover problem for \mathcal{P} in $B \cup R$.*

Proof. As in the proof of Corollary 6, realise G as an intersection graph of blue and red disks. Take for \mathcal{P} one point from each blue–red intersection. The set covers for this instance are exactly the vertex covers of G . ◀

3.1.4 Dominating set of pseudodisks

► **Theorem 9** ([14, 15]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the minimum dominating set problem for pseudodisks.*

► **Corollary 10.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there is a set \mathcal{D} of disks in \mathbb{R}^2 and two dominating sets B and R of \mathcal{D} of at least n disks each such that*

$|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and B is a λ -locally optimal solution to the dominating set problem for \mathcal{D} .

Proof. The instance that was proposed for set cover (Figure 5b) becomes an instance of dominating set when the points of \mathcal{P} are seen as zero-radius disks, i.e. take $\mathcal{D} = \mathcal{P} \cup B \cup R$. A feasible solution that involves some of the zero-radius disks of \mathcal{P} can be transformed into a solution of at most the same cardinality whose support is entirely blue and red since the disks of \mathcal{P} are fully included in the other disks. Thus it suffices to examine the efficiency of local search on blue–red solutions. The blue–red dominating sets of this instance are exactly the covers of points by blue and red disks. ◀

3.1.5 Unit-capacity packing problems

► **Theorem 11** ([12]). *Local search with radius $O(\epsilon^{-2})$ is a $(1 + \epsilon)$ -approximation algorithm for the maximum unit-capacity point-packing problem for disks.*

► **Corollary 12.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two sets B and R of at least n points in \mathbb{R}^2 and a set \mathcal{D} of $\Theta(|R|)$ disks in \mathbb{R}^2 such that every disk of \mathcal{D} contains one point from B and one point from R , $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally optimal solution to the unit-capacity point-packing problem for \mathcal{D} in $B \cup R$.*

Proof. Take for \mathcal{D} the set of all disks associated with the edges, as in the hitting set instance (see Figure 4b). Since every such disk contains only two points of \mathcal{P} , the ‘unit-capacity point-packings’ of this instance are exactly the independent sets of G . The result then follows from the analysis in Corollary 6. ◀

The dual problem, defined in [12], is the unit-capacity disk-packing problem, where we are given a set \mathcal{D} of disks and a set \mathcal{P} of points and we must return a largest subset of \mathcal{D} that covers every point of \mathcal{P} at most once.

► **Corollary 15.** *There is a positive constant C and a positive integer λ_0 such that for every integer $\lambda \geq \lambda_0$ there is a positive integer n_λ such that for every integer $n \geq n_\lambda$ there are two sets B and R of at least n disks each and a set \mathcal{P} of $\Theta(|R|)$ points such that every point of \mathcal{P} is contained in one disk from B and one disk from R , $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and R is a λ -locally maximal solution to the unit-capacity disk-packing problem for $B \cup R$ in \mathcal{P} .*

3.2 Problems with hereditary separators

The paper by Har-Peled and Quanrud [16] is to the best of our knowledge the most extensive study of geometric local search in non-planar settings. The authors study graphs with polynomial expansion⁶, which have strongly sub-linear separators, and in particular intersection graphs of low-density families of objects⁷.

We are still able to give some lower bounds on the local search radii that achieve PTASs. Fix positive integers d , $\lambda \geq \lambda_d$ and n , and let G be the λ -expanding graph built in section 2 on vertex sets B_n and R_n that has $|B_n|, |R_n| = \Theta(n)$ and achieves $|B_n| \geq (1 + c\lambda^{-1/d} - o(1))|R_n|$. Recall that G and its subgraphs have the separator property with $s = 1/d$.

By combining Theorem 3.2.1 and Lemma 2.2.9 from [16], we obtain the following.

⁶ A survey on expansion and sparsity is the book by Nešetřil and Ossona de Mendez [24].

⁷ A family of objects in \mathbb{R}^d has **density** ρ if for any $r \geq 0$ any ball of diameter r intersects at most ρ objects of diameter larger than r and **depth** D if no point of \mathbb{R}^d is contained in $D + 1$ objects.

► **Theorem 16.** *On graphs with hereditary separators of size $O(n^{1-s})$, local search with radius $O(\epsilon^{-s})$ is a $(1 + \epsilon)$ -approximation algorithm for maximum independent set.*

► **Corollary 17.** *For every positive integers d and λ , there are arbitrarily large bipartite graphs on vertex sets (B, R) with hereditary separators of size $O(n^{1-1/d})$ such that $|B| \geq 1 + \Omega(\lambda^{-1/d})|R|$ and R is a λ -locally maximal independent set.*

Proof. Since the graph G is bipartite, both B_n and R_n are independent sets, and by the same analysis as in the proof of Corollary 6 the feasible solution R_n is $(\lambda - 1)$ -locally optimal. ◀

► **Theorem 18** ([16]). *On graphs with hereditary separators of size $O(n^{1-s})$, local search with radius $O(\epsilon^{-O(1)})$ is a $(1 + \epsilon)$ -approximation algorithm for minimum vertex cover.*

► **Corollary 19.** *For every positive integers d and λ , there are arbitrarily large bipartite graphs on vertex sets (B, R) with hereditary separators of size $O(n^{1-1/d})$ such that $|B| \geq 1 + \Omega(\lambda^{-1/d})|R|$ and B is a λ -locally minimal vertex cover.*

Proof. In G both B_n and R_n are vertex covers. Since G is λ -expanding, B_n is λ -locally optimal. ◀

3.3 Matchings and local versions of Hall's theorem

With our terminology, Hall's theorem is as follows.

► **Theorem 20** (Hall's marriage theorem). *Any bipartite graph on vertex sets (B, R) that is $|B|$ -expanding has a matching with $|B|$ edges.*

Restricting the condition to λ -expansion for some fixed λ breaks this property—the matchings of $K_{|B|, \lambda}$ have at most λ edges. However it was observed by Antunes, Mathieu and Mustafa [3] that a strengthening of Hall's theorem holds for planar graphs.

► **Theorem 21.** *There is an absolute constant $c > 0$ such that, for every given integer $\lambda \geq 3$, any bipartite planar graph on vertex sets (B, R) that is λ -expanding has a matching with at least $(1 - c\lambda^{-\frac{1}{2}})|B|$ edges.*

Now it follows from our constructions that this is tight.

► **Corollary 22.** *There are absolute constants $c_0, \lambda_0 > 0$ such that, for every given integer $\lambda \geq \lambda_0$, some bipartite, λ -expanding planar graph on vertex sets (B, R) does not have matchings with more than $(1 - c_0\lambda^{-\frac{1}{2}})|B|$ edges.*

4 Perspectives and open questions

We emphasise that our results apply to standard, non-specialised local-search techniques. Although the approximation quality of a previously successful one-size-fits-all approach cannot be improved, custom algorithms tailored for specific problems can bypass this bound, especially when the exchange graphs are extremely sparse. For example we do not know whether our constructions can be transformed into a local-search-defeating instance for the problem of terrain guarding, a question that can be formulated as follows.

► **Question.** Are the exchange graphs of Gibson et al. [13] for terrain guarding sparser than other planar graphs? What is the minimum size of their separators?

References

- 1 Pankaj K. Agarwal and Nabil H. Mustafa. Independent set of intersection graphs of convex objects in 2D. *Computational Geometry*, 34(2):83–95, 2006. doi:10.1016/j.comgeo.2005.12.001.
- 2 Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society*, 3(4), 10 1990. doi:10.1090/S0894-0347-1990-1065053-0.
- 3 Daniel Antunes, Claire Mathieu, and Nabil H. Mustafa. Combinatorics of local search: An optimal 4-local hall’s theorem for planar graphs. In Kirk Pruhs and Christian Sohler, editors, *25th Annual European Symposium on Algorithms (ESA 2017)*, volume 87 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 8:1–8:13, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ESA.2017.8.
- 4 Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristics for k -median and facility location problems. *SIAM Journal on Computing*, 33(3):544–562, 2004. doi:10.1137/S0097539702416402.
- 5 Rom Aschner, Matthew J. Katz, Gila Morgenstern, and Yelena Yuditsky. Approximation schemes for covering and packing. In Subir Kumar Ghosh and Takeshi Tokuyama, editors, *WALCOM: Algorithms and Computation: 7th International Workshop, WALCOM 2013, Kharagpur, India, February 14-16, 2013. Proceedings*, pages 89–100, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. doi:10.1007/978-3-642-36065-7_10.
- 6 Norbert Bus, Shashwat Garg, Nabil H. Mustafa, and Saurabh Ray. Limits of local search: Quality and efficiency. *Discrete & Computational Geometry*, 57(3):607–624, 4 2017. doi:10.1007/s00454-016-9819-x.
- 7 Sergio Cabello and David Gajser. Simple PTAS’s for families of graphs excluding a minor. *Discrete Applied Mathematics*, 189:41–48, 2015. doi:10.1016/j.dam.2015.03.004.
- 8 Timothy M. Chan and Elyot Grant. Exact algorithms and APX-hardness results for geometric packing and covering problems. *Comput. Geom.*, 47(2, Part A):112–124, 2014. Special Issue: 23rd Canadian Conference on Computational Geometry (CCCG11). doi:10.1016/j.comgeo.2012.04.001.
- 9 Timothy M. Chan and Sariel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry*, 48(2):373–392, 09 2012. doi:10.1007/s00454-012-9417-5.
- 10 Brent N. Clark, Charles J. Colbourn, and David S. Johnson. Unit disk graphs. *Discrete Mathematics*, 86(1):165–177, 1990. doi:10.1016/0012-365X(90)90358-0.
- 11 John Conway and Neil J. A. Sloane. *Sphere Packings, Lattices and Groups*. Grundlehren der mathematischen Wissenschaften. Springer-Verlag New York, New York, NY, USA, 1999. doi:10.1007/978-1-4757-6568-7.
- 12 Alina Ene, Sariel Har-Peled, and Benjamin Raichel. Geometric packing under non-uniform constraints. In *Symposium on Computational Geometry 2012, SoCG ’12, Chapel Hill, NC, USA, June 17-20, 2012*, pages 11–20, 2012. doi:10.1145/2261250.2261253.
- 13 Matt Gibson, Gaurav Kanade, Erik Krohn, and Kasturi Varadarajan. An approximation scheme for terrain guarding. In Irit Dinur, Klaus Jansen, Joseph Naor, and José Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, pages 140–148, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. doi:10.1007/978-3-642-03685-9_11.
- 14 Matt Gibson and Imran A. Pirwani. Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier. In Mark de Berg and Ulrich Meyer, editors, *Proceedings of the 18th*

- Annual European Symposium on Algorithms (ESA)*, pages 243–254, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. doi:10.1007/978-3-642-15775-2_21.
- 15 Sathish Govindarajan, Rajiv Raman, Saurabh Ray, and Aniket Basu Roy. Packing and covering with non-piercing regions. In Piotr Sankowski and Christos Zaroliagis, editors, *Proceedings of the 22nd Annual European Symposium on Algorithms (ESA)*, volume 57 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:17, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ESA.2016.47.
 - 16 Sariel Har-Peled and Kent Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. In *Proceedings of the 23rd Annual European Symposium on Algorithms (ESA)*, pages 717–728, 2015. doi:10.1007/978-3-662-48350-3_60.
 - 17 Sariel Har-Peled and Kent Quanrud. Notes on approximation algorithms for polynomial-expansion and low-density graphs, 2016. arXiv:1603.03098.
 - 18 Dorit S. Hochbaum and Wolfgang Maass. Fast approximation algorithms for a non-convex covering problem. *Journal of Algorithms*, 8(3):305–323, 1987. doi:10.1016/0196-6774(87)90012-5.
 - 19 Lynn H. Loomis and Hassler Whitney. An inequality related to the isoperimetric inequality. *Bulletin of the American Mathematical Society*, 55(10):961–962, 1949. doi:10.1090/S0002-9904-1949-09320-5.
 - 20 Dániel Marx. Efficient approximation schemes for geometric problems? In Gerth Stølting Brodal and Stefano Leonardi, editors, *Algorithms – ESA 2005: 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005. Proceedings*, pages 448–459, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg. doi:10.1007/11561071_41.
 - 21 Dániel Marx. Parameterized complexity of independence and domination on geometric graphs. In Hans L. Bodlaender and Michael A. Langston, editors, *Parameterized and Exact Computation: Second International Workshop, IWPEC 2006, Zürich, Switzerland, September 13-15, 2006. Proceedings*, pages 154–165, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11847250_14.
 - 22 Gary L. Miller, Shang-Hua Teng, William Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. ACM*, 44(1):1–29, 1997. doi:10.1145/256292.256294.
 - 23 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 12 2010. doi:10.1007/s00454-010-9285-9.
 - 24 Jaroslav Nešetřil and Patrice Ossona de Mendez. *Sparsity*, volume 28 of *Algorithms and Combinatorics*. Springer-Verlag Berlin Heidelberg, 2012. doi:10.1007/978-3-642-27875-4.