

A Tight Analysis of Geometric Local Search

Bruno Jartoux · Nabil H. Mustafa

Abstract The last decade has seen the resolution of several basic NP-complete problems in geometric combinatorial optimisation—interestingly, all with the same algorithm: local search. This includes 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. More precisely, 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 $\lambda = \Theta(\varepsilon^{-2})$ yields a PTAS with a running time of $n^{O(\varepsilon^{-2})}$.

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

Our main result is that the approximation guarantee of the standard local search algorithm cannot be improved for any of these problems, which we show by constructing instances with poor “locally optimal solutions”. The key ingredient, of independent interest, is a new lower bound on locally expanding planar graphs. Our construction extends to other graph families with small separators.

Keywords local search · expansion · matchings · Hall’s marriage theorem

Mathematics Subject Classification (2010) 68W25 · 90C27 · 68R10

This work was supported by ANR grant SAGA (JCJC-14-CE25-0016-01) and ANR grant ADDS (ANR-19-CE48-0005). A conference version of this paper was presented at the 34th International Symposium on Computational Geometry (SoCG 2018).

B. Jartoux
Ben-Gurion University of the Negev, Be’er-Sheva, Israel
E-mail: jartoux@post.bgu.ac.il

N. H. Mustafa
Université Gustave Eiffel, LIGM, Equipe A3SI, ESIEE Paris, Noisy-le-Grand, France
E-mail: mustafan@esiee.fr

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 [20, 13] for hardness results). 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.

Hitting set for pseudodisks [26]: given a set \mathcal{P} of points and a family \mathcal{D} of pseudodisks in the plane, compute a smallest subset of \mathcal{P} that intersects all pseudodisks in \mathcal{D} .

Independent set of pseudodisks [1, 12]: given a family \mathcal{D} of pseudodisks in the plane, compute a maximum size subset of pairwise disjoint pseudodisks in \mathcal{D} .

Dominating set of pseudodisks [17, 6]: given a family \mathcal{D} of pseudodisks in the plane, compute a smallest subset of pseudodisks of \mathcal{D} that together intersect all other pseudodisks of \mathcal{D} .

Set cover for disks [11, 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} whose union cover all points of \mathcal{P} .

Constant-capacity point-packing: given a set \mathcal{P} of points, a set \mathcal{D} of disks and a positive integer $k \leq |\mathcal{P}|$, compute a largest subset of \mathcal{P} that hits no disk of \mathcal{D} more than k times [15, 28].

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 maximization problem, e.g., maximum independent set) or decrease (in the case of a minimization 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} for a subset of $X \setminus \mathcal{S}$ of size at least $|\mathcal{S}'| + 1$ (for maximization problems) or at most $|\mathcal{S}'| - 1$ (for minimization 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 $\mathcal{O} \neq \mathcal{S}$ be an optimal solution. We can assume that these solutions are disjoint by considering $\mathcal{S} \setminus \mathcal{O}$ and $\mathcal{O} \setminus \mathcal{S}$ rather than \mathcal{S} and \mathcal{O} : as

$$\frac{|\mathcal{S}|}{|\mathcal{O}|} = \frac{|\mathcal{S} \setminus \mathcal{O}| + |\mathcal{S} \cap \mathcal{O}|}{|\mathcal{O} \setminus \mathcal{S}| + |\mathcal{S} \cap \mathcal{O}|},$$

if we prove (say) that $|\mathcal{S} \setminus \mathcal{O}|/|\mathcal{O} \setminus \mathcal{S}| \leq 1 + \varepsilon$ then this is also true of $|\mathcal{S}|/|\mathcal{O}|$. 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 the following *local vertex expansion property*¹:

$$\mathbf{Minimization:} \text{ for all } \mathcal{S}' \subseteq \mathcal{S} \text{ of size at most } \lambda, |N(\mathcal{S}')| \geq |\mathcal{S}'|. \quad (1)$$

$$\mathbf{Maximization:} \text{ for all } \mathcal{O}' \subseteq \mathcal{O} \text{ of size at most } \lambda, |N(\mathcal{O}')| \geq |\mathcal{O}'|. \quad (2)$$

The construction of exchange graphs exploits the geometric properties of locally optimal solutions and was initially problem-specific. 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} . However, in a recent paper Raman and Ray show that the problems listed above are subcases of a more general problem about the *intersection hypergraph* of two families of pseudodisks [28]. They propose a single algorithm to build exchange graphs for all said problems.

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** if for all $B' \subseteq B$ of size at most λ we have $|N(B')| \geq |B'|$. Note that the roles of B and R in this definition are *not* symmetric. 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 $s = 1$ (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**.

The following result can be pieced together from the analysis of local search by Aschner et al. [5] and is also implicit in previous papers [26, 12] for $s = 1/2$.

Theorem A ([26, 12, 5]). *If a 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 .*

Cabello and Gajser [9] describe a subcase of this theorem for K_h -minor-free graphs, which have separators of size $O(h^{3/2}\sqrt{n})$. Finally, Har-Peled and Quanrud [18, 19] 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 minimization problems, apply Theorem A with $B = \mathcal{S}$ and $R = \mathcal{O}$, and get $|\mathcal{S}| \leq (1 + c_s \lambda^{-s}) \cdot |\mathcal{O}|$. For maximization 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 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 neighbors of the vertices of V' in G .

1.2 Computational efficiency of geometric local search

Given a parameter $\varepsilon > 0$, local search with radius $\lambda = \Theta(\varepsilon^{-\frac{1}{s}})$ provides a $(1 + \varepsilon)$ -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 $n^{O(\varepsilon^{-1/s})}$. In particular, this gives a PTAS in time $n^{O(\varepsilon^{-2})}$ for the five problems listed in section 1.

The parameterized versions of these problems are W[1]-hard: even for unit disks, independent set is W[1]-complete [23], dominating set is W[1]-hard [24], and dominating set of unit disks is easily reduced to our other three problems. The details and reductions are explained in section 5. 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 $\text{poly}(n) \cdot f(\varepsilon)$ for any arbitrary computable function f . In other words, the dependence of the exponent of n on ε is necessary.

Still, the running time of local search is prohibitively expensive as soon as λ is greater than 3 or 4. There have been two complementary approaches towards further progress: firstly, careful implementations that find local improvements more efficiently than by brute force [8]. The second, more structural approach is to better analyze the quality of solutions resulting from local search algorithms, mainly by studying the properties of exchange graphs [3].

2 Our results

2.1 Tightness of Theorem A

Our main result is a construction, given in section 3, showing that Theorem A is asymptotically tight whenever $\frac{1}{s}$ is an integer.

Theorem 1. *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 $(B_n, R_n; E_n)_{n \in \mathbb{N}}$ such that each graph in this family*

- is λ -expanding,
- has the separator property for $s = \frac{1}{d}$, and so does its subgraphs,
- satisfies $|R_n| = n + o(n)$ 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 A restricted to planar graphs are also tight.

2.2 Algorithmic consequences

The analysis of local search in terms of the radius achieving a $(1 + \varepsilon)$ -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 4).

Theorem B ([26]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the minimum hitting set problem for pseudodisks.*

Corollary 1. *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 C ([12]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the maximum independent set problem for pseudodisks.*

Corollary 2. *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 D ([5, 11]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the minimum set cover problem for disks.*

Corollary 3. *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 a set \mathcal{P} of $\Theta(|R|)$ points in \mathbb{R}^2 and two sets B and R of at least n disks in \mathbb{R}^2 , each covering all of \mathcal{P} , and such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and B is a λ -locally optimal solution to the set cover problem for \mathcal{P} in $B \cup R$.*

Theorem E ([17, 6]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the minimum dominating 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 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 F ([15]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the maximum unit-capacity point-packing problem for disks.*

Corollary 5. *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 $1 + \Theta(\lambda^{-1/2})$ ”.

3 Proof of Theorem 1

In this section we build a family of graphs that have the properties stated in Theorem 1. Namely, given parameters d , a large enough λ and n , we construct a bipartite graph G with vertex set (B, R) such that:

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

The vertices of the set R will be called the **red vertices**, and the vertices of the set 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 $x_i(p)$.

Construction of G . 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 \mathcal{E} 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 being a unit d -cube defined by precisely 2^d vertices of \mathcal{E} . In every cell of \mathcal{E} , the vertex with the lexicographically minimum coordinates among the 2^d vertices defining the cell 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 \mathcal{E} .

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 \mathcal{E} . We next place a blue vertex at the center of each of the L^d cells of \mathcal{E} —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 \mathcal{E} , one is connected to all the red vertices of the cell except for $(0, \dots, 0, 1)$ 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 $\tau \in \{0, \dots, t-1\}^d$, where G^τ denotes the translate of $G(d, L)$ by $L \cdot \tau$. The blue vertices of $G(d, L, t)$ are simply the disjoint union of the blue vertices of each G^τ ; the red vertices are also the union of the red vertices of each G^τ , except that we identify duplicate red vertices shared by the boundary of two adjacent grids. See Figure 1 for an example for the case $d = 2$.




Fig. 1: The graph $G(d, L)$ (left, with $d = 2$ and $L = 3$) has L^d grid cells. It is the basic building block of the graph $G(d, L, t)$ (right, with $t = 5$). Red vertices are denoted by squares while the blue vertices are denoted by disks.




Fig. 2: The three-dimensional lowest cell of $G(3, L)$.

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

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

The $L^d + 1$ blue vertices of the translate G^τ form the set B_τ . For the red vertices, note that the outer red vertices of each translate of $G(d, L)$ may be shared between up to 2^d translates. To avoid this overlap, let R_τ consist only of the L^d red vertices $v \in G^\tau$ such that $x_i(v) < L(\tau_i + 1)$ 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^τ .

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

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

where all unions are disjoint. Observe that

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

This concludes the construction of $G(d, L, t)$, although the values of L and t are still to be chosen. We now prove that this graph has the desired expansion and separator properties.

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) neighbors 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 called **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 $|R'| \geq |B'|$. To this end charge each vertex of B' to the anchor red vertex of its cell. For those G^{τ} containing two blue vertices of B' in the lowest cell, one of them remains uncharged. On the other hand, each vertex of R' receives one charge, except the boundary vertices which receive zero charge. Now for each τ for which G^{τ} contains at least two boundary red vertices in R' , charge the uncharged blue vertex of B' in G^{τ} (if it exists) to any one of these (at least two) boundary vertices.

There still remain some uncharged blue vertices, namely one in each of those G^{τ} that have less than two boundary red vertices. However, for each such τ , we show (Lemma 1 below) that G^{τ} must contain at least $\frac{L^2}{2}$ vertices of B' . Thus overall, there are at most $\frac{|B'|}{L^2/2} = \frac{2|B'|}{L^2}$ such τ 's and therefore at most that many uncharged blue vertices of B' . 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 (see the isoperimetric inequality below). By our charging scheme, at least half of them, i.e. $\frac{c_2}{2} \cdot \sqrt{|B'|}$, are still uncharged. Thus when $\frac{2|B'|}{L^2} \leq \frac{c_2}{2} \cdot \sqrt{|B'|}$, or equivalently $|B'| \leq \left(\frac{c_2}{4}\right)^2 \cdot L^4$, the number of uncharged blue vertices will be less than the number of uncharged red vertices. Now setting $L = \Theta\left(\lambda^{\frac{1}{4}}\right)$ and $t = \frac{\sqrt{n}}{L}$ implies that $|R| = n + o(n)$, $|B| = n\left(1 + \frac{1}{L^2}\right) = n\left(1 + \Theta\left(\frac{1}{\lambda^{1/2}}\right)\right)$, and we are done.

Now we present the complete proof for general d . We need two preliminary statements, Lemma 1 and Equation 6. Let the indicator variable η_{τ} be 1 if both blue vertices in the lowest cell of G^{τ} belong to B' and 0 otherwise. Also let δ_{τ} be the number of boundary vertices in R_{τ} . The total number of boundary vertices in R' is thus

$$\delta = |R_b \cap R'| + \sum_{\tau} \delta_{\tau}. \quad (4)$$

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

Proof. As B' contains both blue vertices from the lowest cell of G^τ , R' contains the 2^d red vertices of this cell. If $\delta_\tau = 0$, that is, R_τ contains no boundary vertex, then the blue vertex in each of the other cells of G^τ is present in B' , and so B' includes all of B_τ , which consists of $L^d + 1$ blue vertices.

It remains to consider the case when R_τ contains a unique boundary vertex $v_r \in R' \cap R_\tau$. Without loss of generality, assume that $\tau = (0, \dots, 0)$. As both blue vertices from the lowest cell of G^τ belong to B' , the boundary vertex v_r cannot be the lowest vertex of G^τ , 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^τ , 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^τ 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 another boundary red vertex in one of these L cells.

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

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

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

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, $i \in [d]$, 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. As such, it contains at least one boundary vertex: among its vertices with maximum i -th coordinate, the lexicographically smallest is either in R_b or the anchor vertex of an empty cell. Thus the total number δ of boundary vertices is at least $|S_i|$. The combinatorial Loomis–Whitney inequality [22] 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

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

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

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

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

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

with the last equality following from Equation (4). Use the lower bounds Equation (5) and Equation (6) for the second and third summands to get

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

Now $|R'| \geq |B'|$ 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 $2^{3d-1}|B'| \leq L^{d^2}$. \square

Setting parameters. Given d, λ and n , choose

$$L = \max \left\{ 2, \left\lceil \left(2^{3d-1} \lambda \right)^{1/d^2} \right\rceil \right\} \text{ and } t = \lceil n^{1/d} L^{-1} \rceil.$$

Note that L does not depend on n and L^d is $\Theta(\lambda^{1/d})$ when $\lambda \rightarrow +\infty$. Using Equation (3), we obtain the required bounds on R and B :

$$\begin{aligned} |R| &= (tL+1)^d = n + o(n), \\ \frac{|B|}{|R|} &= \frac{t^d (L^d+1)}{(tL+1)^d} = 1 + \frac{1}{L^d} + o(1) \text{ as } n \rightarrow +\infty \\ &\geq 1 + c_d \lambda^{-\frac{1}{d}} + o_n(1) \text{ for } \lambda \geq \lambda_d, \end{aligned}$$

where the positive constants c_d and λ_d depend only on d . Since $2^{1-3d}L^{d^2} \geq \lambda$, it follows from Lemma 2 that property (2) holds: G is λ -expanding. This achieves the proof of Theorem 1.

Finally, the stated separator properties are a consequence of the specific geometric structure of $G(d, L, t)$.




Fig. 4: A tight instance for the hitting set problem.

4.1.1 Hitting set for pseudodisks

Theorem B ([26]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the minimum hitting set problem for pseudodisks.*

Corollary 1. *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 neighbors, and thus B is λ -locally optimal. \square

4.1.2 Independent set of pseudodisks

Theorem C ([12]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the maximum independent set problem for pseudodisks.*

Corollary 2. *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 [21].) 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 maximization 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. \square

4.1.3 Set cover for disks




Fig. 5: “Tight” instances for independent set and set cover with disks.

Theorem D ([5, 11]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the minimum set cover problem for disks.*

Corollary 3. *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 a set \mathcal{P} of $\Theta(|R|)$ points in \mathbb{R}^2 and two sets B and R of at least n disks in \mathbb{R}^2 , each covering all of \mathcal{P} , and such that $|B| \geq (1 + C\lambda^{-\frac{1}{2}})|R|$ and B 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 2, 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 . \square

4.1.4 Dominating set of pseudodisks

Theorem E ([17, 6]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the minimum dominating 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 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. \square

4.1.5 Constant-capacity packing problems

Our results hold already for the unit-capacity version of these problems.

Theorem F ([15]). *Local search with radius $O(\varepsilon^{-2})$ is a $(1 + \varepsilon)$ -approximation algorithm for the maximum unit-capacity point-packing problem for disks.*

Corollary 5. *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 2. \square

The dual problem, defined in [15], 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 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 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} .*

4.2 Problems with hereditary separators

The paper by Har-Peled and Quanrud [18] 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 expansion, which have strongly sub-linear separators. A survey on expansion and sparsity is the book by Nešetřil and Ossona de Mendez [27].

In particular, this includes intersection graphs of low-density families of objects. (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.)

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 3 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 [18], we obtain the following.

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

Corollary 7. *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 2 the feasible solution R_n is $(\lambda - 1)$ -locally optimal. \square

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

Corollary 8. *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. \square

4.3 Matchings and local versions of Hall's theorem

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

Theorem I (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—e.g., any matching in $K_{|B|, \lambda}$ has at most λ edges. However it was observed by Antunes et al. [3] that a strengthening of Hall's theorem holds for planar graphs.

Theorem J. *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 9. *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.*

5 Parameterized hardness and conditional lower bounds

Lower bounds on the possible running time of PTASs for our problems are derived from results in parameterized complexity in the manner of [23].

Every combinatorial optimization problem has a **decision version** parameterized by $k \in \mathbb{N}$, informally “Does the instance have solutions smaller than k (minimization) / larger than k (maximization)?”.

If this can be answered in time $f(k)n^c$ for some arbitrary computable function f and $c \geq 0$, then (the decision version of) the problem is **fixed-parameter tractable**, or in the class FPT. Parameterized problems that are reducible to the search for a clique of size k in general graphs—under an appropriate notion of reduction—form the class W[1]. The inclusion $\text{FPT} \subseteq \text{W}[1]$ is thought to be proper; this is for example a consequence of the exponential time hypothesis.

Theorem K ([23]). *Maximum independent set of unit disks is W[1]-complete.*

Theorem L ([24]). *Minimum dominating set of unit disks is W[1]-hard.*

This extends to hitting set by a simple reduction.

Corollary 10. *Minimum hitting set of unit disks is W[1]-hard.*

Proof. Given an instance of dominating set with n unit disks, let \mathcal{P} consist of their centers and \mathcal{D} consist of the n radius-2 disks centered at \mathcal{P} . The hitting sets of \mathcal{D} in \mathcal{P} are exactly the dominating sets of the initial unit disks. \square

Corollary 11. *Minimum set cover by unit disks is W[1]-hard.*

Proof. Given an instance of hitting set with unit disks \mathcal{D} and points \mathcal{P} , let \mathcal{P}' consist of the centers of \mathcal{D} and \mathcal{D}' contain one unit disk centered at each point in \mathcal{P} . The set covers of \mathcal{P}' in \mathcal{D}' are exactly the hitting sets of \mathcal{D} in \mathcal{P} . \square

Recall the definition of the unit-capacity disk-packing problem on page 14.

Corollary 12. *Unit-capacity disk-packing is W[1]-hard, even with unit disks.*

Proof. Given an instance of independent set with n unit disks, let \mathcal{D} consist of these disks and \mathcal{P} contain one point from each nonempty pairwise intersection in \mathcal{D} . The set \mathcal{P} has size $O(n^2)$ and can be computed in $O(n^2)$. Now the feasible solutions of unit-capacity disk-packing for \mathcal{D} and \mathcal{P} are exactly the independent sets of \mathcal{D} . \square

Finally the disk-packing problem can be reduced to the point-packing one (defined on page 2).

Corollary 13. *Unit-capacity point-packing is W[1]-hard, even with unit disks.*

Proof. Given an instance of unit-capacity disk-packing with unit disks \mathcal{D} and points \mathcal{P} , let \mathcal{D}' consist one unit disk centered at each point in \mathcal{P} and \mathcal{P}' contain the centers of the disks in \mathcal{D} . The feasible solutions of unit-capacity disk-packing for \mathcal{D} and \mathcal{P} are exactly the feasible solutions of unit-capacity point-packing for \mathcal{D}' and \mathcal{P}' . \square

An **efficient PTAS** is a PTAS that runs in time $f(\varepsilon)n^c$ for an arbitrary function f and $c \geq 0$. The following observation first appeared in C. Bazgan’s master’s thesis [7, 10].

Theorem M. *A combinatorial optimization problem with an integral-valued objective function and an efficient PTAS with f computable is in FPT when parameterized by the value of its solution.*

Proof. Given a minimization problem with an efficient PTAS whose running time is $f(\varepsilon)n^c$, for each instance we can find in time $f(0.5/k)n^c$ a solution with value V such that

$$\left(1 + \frac{1}{2k}\right) \cdot \text{OPT} \geq V \geq \text{OPT}.$$

If $V \geq k + 1$, then

$$\text{OPT} \geq \frac{2kV}{2k+1} \geq \frac{2k^2 + 2k}{2k+1} > k,$$

so computing V answers the question “Does the instance have $\text{OPT} \leq k$?” in time $f(0.5/k)n^c$, i.e. the problem parameterized by the value of its solution is in FPT. \square

Since the problems that we consider are W[1]-hard, we conclude that none of them admit PTASs with running time $f(\varepsilon)n^{O(1)}$ for any computable f unless the inclusion $\text{FPT} \subseteq \text{W}[1]$ is an equality.

6 Perspectives and open questions

We emphasize that our results apply to standard, non-specialized 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. [16] for terrain guarding sparser than other planar graphs? What is the minimum size of their separators?

Acknowledgements We thank the referees of previous versions of this article for several helpful comments.

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. Agarwal, P.K., Mustafa, N.H.: 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. Alon, N., Seymour, P., Thomas, R.: A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society* **3**(4), 801–808 (1990). DOI 10.1090/S0894--1990-1065053-0
3. Antunes, D., Mathieu, C., Mustafa, N.H.: Combinatorics of local search: An optimal 4-local hall’s theorem for planar graphs. In: K. Pruhs, C. Sohler (eds.) 25th Annual European Symposium on Algorithms (ESA 2017), *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 87, pp. 8:1–8:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2017). DOI 10.4230/LIPIcs.ESA.2017.8
4. Arya, V., Garg, N., Khandekar, R., Meyerson, A., Munagala, K., Pandit, V.: Local search heuristics for k -median and facility location problems. *SIAM Journal on Computing* **33**(3), 544–562 (2004). DOI 10.1137/S0097539702416402
5. Aschner, R., Katz, M.J., Morgenstern, G., Yuditsky, Y.: Approximation schemes for covering and packing. In: S.K. Ghosh, T. Tokuyama (eds.) WALCOM: Algorithms and Computation: 7th International Workshop, WALCOM 2013, Kharagpur, India, February 14–16, 2013. Proceedings, pp. 89–100. Springer Berlin Heidelberg, Berlin, Heidelberg (2013). DOI 10.1007/978-3-642-36065-7_10
6. Basu Roy, A., Govindarajan, S., Raman, R., Ray, S.: Packing and covering with non-piercing regions. In: P. Sankowski, C. Zaroliagis (eds.) Proceedings of the 22nd Annual European Symposium on Algorithms (ESA), *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 57, pp. 47:1–47:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2016). DOI 10.4230/LIPIcs.ESA.2016.47
7. Bazgan, C.: Schémas d’approximation et complexité paramétrée. DEA memoir, Université Paris Sud (1995). Unobtainable online
8. Bus, N., Garg, S., Mustafa, N.H., Ray, S.: Limits of local search: Quality and efficiency. *Discrete & Computational Geometry* **57**(3), 607–624 (2017). DOI 10.1007/s00454-016-9819-x
9. Cabello, S., Gajser, D.: 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
10. Cesati, M., Trevisan, L.: On the efficiency of polynomial time approximation schemes. *Information Processing Letters* **64**(4), 165 – 171 (1997). DOI 10.1016/S0020-0190(97)00164-6
11. Chan, T.M., Grant, E.: Exact algorithms and APX-hardness results for geometric packing and covering problems. *Comput. Geom.* **47**(2, Part A), 112–124 (2014). DOI 10.1016/j.comgeo.2012.04.001. Special Issue: 23rd Canadian Conference on Computational Geometry (CCCG11)
12. Chan, T.M., Har-Peled, S.: Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry* **48**(2), 373–392 (2012). DOI 10.1007/s00454-012-9417-5
13. Clark, B.N., Colbourn, C.J., Johnson, D.S.: Unit disk graphs. *Discrete Mathematics* **86**(1), 165 – 177 (1990). DOI 10.1016/0012-365X(90)90358-O
14. Conway, J., Sloane, N.J.A.: 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
15. Ene, A., Har-Peled, S., Raichel, B.: Geometric packing under non-uniform constraints. In: Proceedings of the Twenty-eighth Annual Symposium on Computational Geometry, SoCG ’12, pp. 11–20. ACM, New York, NY, USA (2012). DOI 10.1145/2261250.2261253. URL <http://doi.acm.org/10.1145/2261250.2261253>
16. Gibson, M., Kanade, G., Krohn, E., Varadarajan, K.: An approximation scheme for terrain guarding. In: I. Dinur, K. Jansen, J. Naor, J. Rolim (eds.) 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, pp. 140–148. Springer Berlin Heidelberg, Berlin, Heidelberg (2009). DOI 10.1007/978-3-642-03685-9_11

17. Gibson, M., Pirwani, I.A.: Algorithms for dominating set in disk graphs: Breaking the $\log n$ barrier. In: M. de Berg, U. Meyer (eds.) Proceedings of the 18th Annual European Symposium on Algorithms (ESA), pp. 243–254. Springer Berlin Heidelberg, Berlin, Heidelberg (2010). DOI 10.1007/978-3-642-15775-2_21
18. Har-Peled, S., Quanrud, K.: Approximation algorithms for polynomial-expansion and low-density graphs. In: Proceedings of the 23rd Annual European Symposium on Algorithms (ESA), pp. 717–728. Springer, Berlin, Heidelberg, Berlin, Heidelberg (2015). DOI 10.1007/978-3-662-48350-3_60
19. Har-Peled, S., Quanrud, K.: Notes on approximation algorithms for polynomial-expansion and low-density graphs (2016)
20. Hochbaum, D.S., Maass, W.: Fast approximation algorithms for a nonconvex covering problem. *Journal of Algorithms* **8**(3), 305 – 323 (1987). DOI 10.1016/0196-6774(87)90012-5
21. Koebe, P.: Kontaktprobleme der konformen abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig, Mathematisch-Physische Klasse* **88**, 141–164 (1936)
22. Loomis, L.H., Whitney, H.: 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
23. Marx, D.: Efficient approximation schemes for geometric problems? In: G.S. Brodal, S. Leonardi (eds.) Algorithms – ESA 2005: 13th Annual European Symposium, Palma de Mallorca, Spain, October 3-6, 2005. Proceedings, pp. 448–459. Springer Berlin Heidelberg, Berlin, Heidelberg (2005). DOI 10.1007/11561071_41
24. Marx, D.: Parameterized complexity of independence and domination on geometric graphs. In: H.L. Bodlaender, M.A. Langston (eds.) Parameterized and Exact Computation: Second International Workshop, IWPEC 2006, Zürich, Switzerland, September 13-15, 2006. Proceedings, pp. 154–165. Springer Berlin Heidelberg, Berlin, Heidelberg (2006). DOI 10.1007/11847250_14
25. Miller, G.L., Teng, S.H., Thurston, W., Vavasis, S.A.: Separators for sphere-packings and nearest neighbor graphs. *J. ACM* **44**(1), 1–29 (1997). DOI 10.1145/256292.256294
26. Mustafa, N.H., Ray, S.: Improved results on geometric hitting set problems. *Discrete & Computational Geometry* **44**(4), 883–895 (2010). DOI 10.1007/s00454-010-9285-9
27. Nešetřil, J., Ossona de Mendez, P.: Sparsity, *Algorithms and Combinatorics*, vol. 28. Springer-Verlag Berlin Heidelberg, Berlin, Heidelberg (2012). DOI 10.1007/978-3-642-27875-4
28. Raman, R., Ray, S.: Planar Support for Non-piercing Regions and Applications. In: Y. Azar, H. Bast, G. Herman (eds.) 26th Annual European Symposium on Algorithms (ESA 2018), *Leibniz International Proceedings in Informatics (LIPIcs)*, vol. 112, pp. 69:1–69:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2018). DOI 10.4230/LIPIcs.ESA.2018.69. URL <http://drops.dagstuhl.de/opus/volltexte/2018/9532>