Geometric Optimization Parameterized by Piercing Complexity

Aritra Banik
National Institute of Science Education and Research, Bhubaneswar, India
Homi Bhabha National Institute, Training School Complex, Anushakti Nagar, Mumbai, India

Rajiv Raman
IIIT-Delhi, India

Saurabh Ray
NYU Abu Dhabi, United Arab Emirates

Abstract

Packing and Covering problems with geometric regions in the plane have been extensively studied and several notions of "complexity" of the regions involved have been developed and exploited to obtain good approximation algorithms. Examples of such complexity measures are VC-dimension, union complexity, shallow-cell complexity, fatness, etc. While these restrictions lead to constant-factor approximation algorithms in many cases, they typically do not lead to PTASs. In fact, several geometric Set Cover and Discrete Independent Set variants remain APX-hard even when these parameters are small, as demonstrated in earlier work by Chan and Grant (Exact algorithms and APX-hardness results for geometric packing and covering problems. Comput. Geom., 2014), and by Har-Peled and Quanrud (Approximation Algorithms for Polynomial-Expansion and Low-Density Graphs. SIAM J. Comput., 2017). A key feature of these hardness constructions is that many pairs of regions in the input *pierce* one another.

Motivated by this observation, we initiate a systematic study of geometric families parameterized by their *piercing complexity*. A connected region A is said to pierce a connected region B if $B \setminus A$ has more than one connected component; we consider instances in which every region is pierced by at most a constant number of others. This framework smoothly interpolates between the classical non-piercing case-where local-search PTASs are known due to Raman and Ray (Constructing Planar Support for Non-Piercing Regions, Discret. Comput. Geom., 2020), and the fully general case, where APX-hardness persists.

Our main contribution is to show that bounded-piercing families admit efficient approximation schemes for fundamental geometric optimization problems. For regions in the plane with a constant piercing bound, we obtain PTASs for the (unweighted) *Discrete Independent Set* and *Set Cover* problems, and constant-factor approximation algorithms for their weighted variants. These results strictly generalize the known PTASs for non-piercing families and yield improved guarantees for several long-standing special cases, including Independent Set and Set Cover with axis-parallel rectangles under bounded piercing.

Overall, our work identifies piercing complexity as a robust and expressive topological parameter-distinct from geometric notions such as density or fatness-and demonstrates that bounding this parameter yields a broad family of geometric instances for which PTASs become achievable.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Geometric set cover, geometric discrete independent set, approximation algorithms, PTAS, parameterized complexity, piercing complexity, non-piercing regions, independent set, axis-parallel rectangles

Digital Object Identifier 10.4230/LIPIcs.ICALP.2026.21

Category Track A: Algorithms, Complexity and Games

1 Introduction

Given a set X of n elements and a collection $\mathcal{S}$ of subsets of X whose union covers X, the *Set Cover* problem asks for a smallest sub-collection $\mathcal{S}'$ that also covers X. In the *Discrete Independent Set* problem, the objective is to select a largest sub-collection $\mathcal{S}' \subseteq \mathcal{S}$ such that each element of X is contained in at most one set in $\mathcal{S}'$. Both problems have been extensively studied, and tight bounds on their approximability are known in the general setting [15, 29, 25].

Improved approximation algorithms can often be obtained when these problems are restricted to geometric settings. Such improvements are typically enabled by bounding an appropriate notion of *complexity* of the underlying geometric objects. Well-studied complexity measures include dimension, VC-dimension, union complexity (and its generalization, shallow-cell complexity), and fatness, among others [32, 30, 20, 8]. In particular, small VC-dimension or shallow-cell complexity implies the existence of small ϵ-nets [32], which in turn yield approximation algorithms via LP-rounding.

A representative example is the Set Cover problem defined by unit disks and points in the plane, which admits a PTAS [5]. In contrast, when disks are replaced by geometrically similar regions – such as nearly congruent, nearly circular ellipses – the problem becomes APX-hard, despite remaining small with respect to the above complexity measures [7]. Chan and Grant [7] gave several further examples of simple geometric objects, including axis-parallel strips and wedges, for which both the Set Cover and Discrete Independent Set problems are APX-hard (see also [21]). A key structural feature underlying these hardness constructions is that many pairs of regions in the input *pierce* one another. Informally, a connected region A is said to pierce another connected region B if $B \setminus A$ has two or more connected components.

In contrast, PTASs are known for a wide range of geometric packing and covering problems. In many cases, these algorithms are based on a simple local-search framework [9, 33, 3, 18, 5, 35]. A common requirement in these settings is that the regions are *non-piercing*[1], meaning that no region pierces another (see Definition 2.1).

The fact that there exist PTASs for non-piercing regions, while known APX-hardness reductions involve many piercing pairs of regions motivates the central question of this paper: can one quantify the *amount* of piercing in an instance, and thereby interpolate between the tractable non-piercing regime and the intractable general case?

Quantifying the amount of piercing. We quantify piercing using two parameters. For two regions A and B, we say that B splits A into k pieces if $A \setminus B$ has k connected components. The *fragmentation number* f is the maximum number of pieces into which any region can be split by another, and the *piercing number* r is the maximum number of regions that pierce any single region. A collection of regions with fragmentation number at most f and piercing number at most r is called an (f, r)-piercing set. The two parameters play different roles: r bounds how many topologically "disruptive" interactions any one region has, while f bounds the damage caused by one such interaction. This is a deliberately modest parameterization. For many simple geometric families, such as axis-parallel rectangles, the fragmentation number is already constant, so the piercing number is the parameter that interpolates between the non-piercing regime and highly piercing instances. We do not claim that these are the only possible parameters, but they are sufficient to recover

[1] Examples of non-piercing regions include axis-parallel squares, disks, axis-parallel unit-height rectangles, and homothets of convex regions.

local-search PTASs and they isolate a structural feature absent from standard measures such as VC-dimension or union complexity. This framework interpolates smoothly between the non-piercing case ($r = 0$) and the fully general setting. Our main result, stated informally below, shows that bounded piercing suffices to recover PTASs for fundamental geometric optimization problems. Note that even if f and r are poly-logarithmic in n, our results imply QPTASs for the two problems considered in this paper.

▶ **Theorem** (Informal). *Let Γ be a set of n regions forming an (f, r)-piercing set. For any $\epsilon > 0$, there is an algorithm with running time $O(n^{poly(f,r,1/\epsilon)})$ that yields a $(1 + \epsilon)$-approximation for the Set Cover problem defined by Γ and a point set P. Analogously, there is an algorithm with the same running time that yields a $(1 - \epsilon)$-approximation for the Discrete Independent Set problem.*

Applications. Many geometric optimization problems employ idealized shape models: sensor coverage as disks or wedges [39], map labels as rectangles [2], and facility service regions defined by ℓ_p-norm balls [19]. These models enable elegant algorithms but poorly model realistic scenarios. Non-piercing regions [35] partially address this by allowing highly irregular and non-convex geometric shapes, but the non-piercing requirement is often violated: obstacles may cause sensor regions to pierce one another, labels for intersecting roads can create piercing pairs, and competing service territories may fragment each other. The (f, r)-piercing framework should therefore be viewed as a structural, parameterized relaxation of the non-piercing assumption. It applies to instances in which the piercing graph has bounded maximum degree and each piercing has bounded fragmentation; such restrictions can arise from placement rules, design constraints, or simply from the input being only mildly non-piercing. We use these scenarios as motivation rather than as a claim that the piercing number is universally small in practice.

Improved algorithms for special cases. Our results imply improved approximation guarantees for several well-studied geometric problems. Both the Set Cover and Discrete Independent Set problems are APX-hard for axis-parallel rectangles [7]. For the continuous Independent Set problem, polynomial-time constant-factor approximations [31, 17] and a QPTAS [1] are known, while for the Set Cover problem, the best known approximation is $O(\log \text{OPT})$ via ϵ-nets [6, 14], which is tight under this framework [34]. We show that for both problems, if the piercing number of the rectangles is constant, then a PTAS exists – something not achievable by existing techniques.

A similar phenomenon arises for points and line segments. LP-based methods yield at best an $O(\log \text{OPT})$-approximation [6, 14], and lower bounds on ϵ-nets rule out substantially better guarantees [4]. At the same time, the problem remains APX-hard even for axis-parallel segments [27]. Our results again yield PTASs for both the Set Cover and Discrete Independent Set problems when the piercing number is bounded.

Comparison with other parameters. Parameterized algorithm design based on structural measures of the input is an active research direction [12]. Piercing complexity differs fundamentally from existing parameters in that it is purely topological, unlike geometric parameters such as fatness or density, or combinatorial parameters such as VC-dimension, shallow-cell complexity, or treewidth.

Treewidth, even when combined with bounded ply, can be useful for geometric approximation algorithms [28]. However, intersection graphs of (f, r)-piercing regions can contain arbitrarily large grid minors even with bounded ply, implying unbounded treewidth. Thus, treewidth and piercing complexity are incomparable.

Har-Peled and Quanrud [22] introduced the notion of low-density intersection graphs, which leads to PTASs for several geometric problems. Their framework, however, relies on a geometric notion of object size, whereas (f, r)-piercing is entirely topological, making the two notions incomparable.

Finally, while (f, r)-piercing regions may have arbitrarily large union complexity, we show that they nevertheless have shallow-cell complexity linear in f and r (Theorem 3.3). Since shallow-cell complexity alone is insufficient to yield PTASs [22, 7], this demonstrates that piercing complexity captures a distinct structural property that enables efficient approximation schemes.

Technical contribution. Our algorithms follow the local-search framework used in earlier work on geometric packing and covering problems. For non-piercing regions, prior work constructs planar supports, and the needed separators follow directly from planarity [35]. With bounded piercing, the corresponding exchange graphs need not be planar. The main technical step is to recover enough sparsity for local search: lens-bypassing operations simplify the arrangement while preserving the relevant coverage conditions and the (f, r)-piercing property, and the resulting exchange graph is represented as a bounded-depth intersection graph of connected subgraphs of the planar dual arrangement graph. A Frederickson-style decomposition of this graph then supplies the local pieces used in the PTAS analysis. This is the point at which the proof extends, rather than merely reuses, the non-piercing machinery.

1.1 Related Work

For packing problems with *fat objects*[2] in bounded dimension, eg. balls or cubes, Erlebach et al. [13] obtained a PTAS building on the work of Hochbaum and Maass [24] for unit disks. For regions that are not fat, Adamaszeck and Wiese [1] obtained a QPTAS (including weighted settings) for regions in the plane under the restriction that each region in the set is path connected. In the discrete setting, Chan and Grant [7] obtained hardness results for packing and covering problems with simple geometric regions in the plane such as axis-parallel strips.

For covering problems, approximation algorithms primarily build on ϵ-nets. Haussler and Welzl [23] showed ϵ-nets of size $O(d/\epsilon \log 1/\epsilon)$ for set systems with bounded VC-dimension, leading to $O(\log \textsc{Opt})$-approximations by Brönnimann and Goodrich [6] and Even et al.[14]. Clarkson and Varadarajan [11] improved ϵ-nets for geometric systems, culminating in Varadarajan's [38] *quasi-uniform sampling* which was refined by Chan et al. [8] for systems with shallow-cell complexity $\phi(n)$, yielding ϵ-nets of size $O(\log \phi(n))$ and an $O(\log \phi(n))$-approximation for weighted Set Cover, though with large constants even in simple cases. These algorithms also work in the weighted setting.

For *unweighted* problems, the unifying tool has been the *local search framework*, which iteratively improves a feasible solution with small changes. This yielded PTASs for independent set [9], hitting set [33], and Set Cover and dominating set [5]. Raman and Ray [35] showed that set systems from *non-piercing regions* admit planar supports, implying PTAS results for many packing and covering problems via local search, and also for demand/capacitated versions [36]. Raman and Singh [37] extended these results to higher-genus surfaces. The limitations of the local search approach were studied by Mustafa and Jartoux [26].

[2] There are many definitions of fatness, but a definition that is sufficient for us is that region is fat if the ratio of the inscribing and circumscribing balls have radii that are at most a constant factor apart.

2 Preliminaries

In this section, we introduce notation and define terms that we use throughout the paper. We use the term "region" to refer to a closed and bounded set of points in the plane whose boundary is a simple Jordan curve. Thus, each region in our setting is a simply connected region[3]. We denote the boundary of any region α by $\partial\alpha$.

▶ **Definition 2.1** (Non-piercing set [35]). *A set $\mathcal{R}$ of regions in the plane is said to be* non-piercing *if for any $\alpha, \beta \in \mathcal{R}$, $\alpha \setminus \beta$ is a connected set.*

The families of regions we consider in this paper may contain *piercing* pairs, where for two regions α, β, the set $\alpha \setminus \beta$ consists of two or more connected components. We define below the *fragmentation number* which denotes the number of components in $\alpha \setminus \beta$.

▶ **Definition 2.2** (Fragmentation number). *Let α and β be two regions. If their boundaries intersect, we define the fragmentation number of the ordered pair (α, β), denoted $\textsc{frag}(\alpha, \beta)$, as the number of connected components in $\alpha \setminus \beta$ (see Figure 1). If the boundaries of α and β do not intersect (which can happen if either one is contained in the other or they are disjoint), we define $\textsc{frag}(\alpha, \beta) = \textsc{frag}(\beta, \alpha) = 0$.*

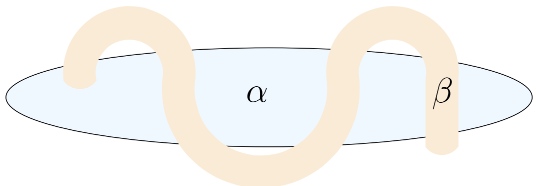

Figure 1 $\textsc{frag}(\alpha, \beta) = 4$.

▶ **Definition 2.3** (Piercing). *We say that two regions α and β are* piercing *if $\textsc{frag}(\alpha, \beta) \geq 2$ and* non-piercing *otherwise.*

▶ **Definition 2.4** (Dual arrangement graph). *Let Γ be an arrangement of regions in the plane. The* dual arrangement graph *$\textsc{dual}(\Gamma)$ of Γ is a graph in which there is a vertex corresponding to each cell in the arrangement of Γ and two vertices are adjacent iff the boundaries of the corresponding cell touch in an arc of non-zero length. Observe that $\textsc{dual}(\Gamma)$ is a plane multi-graph.*

▶ **Definition 2.5** ((f, r)-piercing regions). *We say that a set of regions Γ is (f, r)-piercing if i) for any two regions $\alpha, \beta \in \Gamma$, $\textsc{frag}(\alpha, \beta) \leq f$ and ii) for any region $\alpha \in \Gamma$, the number of regions $\beta \in \Gamma$ that pierce α is at most r.*

▶ **Definition 2.6** (Lens [5, 35]). *Let α and β be any two regions. A* lens *formed by α and β is a cell in the arrangement of the two regions that is contained in both and has two sides, one along the boundary of α and the other along the boundary of β (see Figure 2).*

▶ **Definition 2.7** (Minimal lens [5, 35]). *Let Γ be any set of regions. We say that lens L defined by two of the regions $\alpha, \beta \in \Gamma$ is a* minimal lens *if L does not contain any lens L' defined by any other pair of regions $(\alpha', \beta') \neq (\alpha, \beta)$ in Γ.*

[3] Raman and Ray [35] considered a slightly more general notion of regions, but for our purposes, we work with only simply connected regions

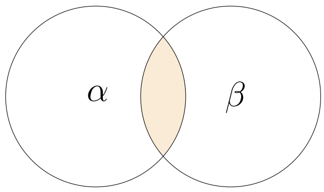

Figure 2 Lens formed by α and β.

▶ **Definition 2.8** (Expansion [5, 35]). *Let Γ be any collection of regions. For any region R (not necessarily in Γ), we define an "expansion of R w.r.t. Γ" as the region $R_+^\Gamma = R \oplus B_\epsilon$ where $\epsilon > 0$, B_ϵ denotes an open ball of radius ϵ around the origin, and $\oplus$ denotes Minkowski sum. Here ϵ is chosen to be any appropriately small quantity so that R_+^Γ does not contain any vertices of the arrangement of Γ other than those lying in R. When Γ is clear from the context, we write R_+ instead of R_+^Γ.*

▶ **Definition 2.9** (Lens bypassing [5, 35]). *Let Γ be any collection of regions. Let L be a lens formed by the regions α and β in Γ. We define the operation "bypassing of L by α" as the modification of α to $\alpha \setminus L_+$.*

The following definition is inspired by the decomposition theorem for planar graphs due to Frederickson [16].

▶ **Definition 2.10.** *We say that a family of graphs $\mathcal{G}$ admits a $(\lambda_1, \lambda_2, \lambda_3)$-decomposition for parameters $\lambda_1, \lambda_2, \lambda_3$ if for any graph $G = (V, E)$ in $\mathcal{G}$ with m vertices, and any given parameter t, we can find a set $X \subseteq V$ of size at most $\lambda_1 m/\sqrt{t}$, and a partition of $V \setminus X$ into $\ell \leq \lambda_2 m/t$ sets $V_1, \ldots, V_\ell$ s.t.*
(i) *$|V_i| \leq t$ for all i*
(ii) *$|N_G(V_i) \cap X| \leq \lambda_3\sqrt{t}$ for all i.*
(iii) *$N_G(V_i) \cap V_j = \emptyset$ for $i \neq j$*
Here, for any $U \subseteq V$, $N_G(U)$ denotes the set of all neighbors of the vertices in U in G.

Frederickson [16] showed that the family of planar graphs admits a (c_1, c_2, c_3)-decomposition for some constants c_1, c_2, c_3.

▶ **Definition 2.11** (Shallow cell complexity of a set of regions [8]). *The shallow cell complexity of a set of regions Γ is a function ϕ s.t. for any k, the number of combinatorially distinct cells of depth $\leq k$ is $O(n \cdot \phi(n) \cdot poly(k))$. Two cells are considered combinatorially equivalent if they are contained in the same set of regions.*

Shallow cell complexity is defined more generally in [8] for set systems. For our purposes, the above equivalent definition suffices.

3 Main Results

We now state the formal algorithmic consequences of bounded piercing complexity. Throughout this section, Γ is an (f, r)-piercing set of regions, where f and r are as in Definition 2.5. Our main contribution is a pair of local-search PTASs for the two unweighted problems considered in the paper. In the *Discrete Independent Set* problem, the objective is to select a largest subset of Γ such that any point in P is contained in at most one selected region. In the *Set Cover* problem, the goal is to select a smallest subset of regions such that each point of P is in at least one selected region. If f and r are constants, the algorithms run in polynomial time; if f and r are polylogarithmic in n, the same bounds give QPTASs.

The proofs use the same broad local-search paradigm as the known PTASs for non-piercing regions, but the structural step is different. Instead of obtaining a planar support, we use the lens-bypassing and decomposition lemmas in Section 4 to build exchange graphs that admit the decompositions required by the local-search analysis. The t-local search algorithm referred to in the following two theorems is described in Section 5.

▶ **Theorem 3.1** (Discrete Independent Set). *The t-local search algorithm with $t = Cr^7 f^2/\epsilon^2$ for an appropriately large C and $\epsilon \in (0,1]$ yields a $(1-\epsilon)$-approximation for the Discrete Independent Set problem defined by a set of (f,r)-piercing regions Γ and a set of points P. The algorithm runs in time $O(n^{O(t)})$.*

▶ **Theorem 3.2** (Set Cover). *The t-local search algorithm with $t = Cr^7 f^2/\epsilon^2$ for an appropriately large C and $\epsilon \in (0,1]$ yields a $(1+\epsilon)$-approximation for the Set Cover problem defined by a set of (f,r)-piercing regions Γ and a set of points P. The algorithm runs in time $n^{O(t)}$.*

We also show that the shallow cell complexity of a set of regions with the above parameters f and r is $O(r^2 f)$ from which the following approximation algorithms for the weighted setting follow via existing results of Chan and Har-Peled [9], and Chan et al. [8].

▶ **Theorem 3.3.** *Let Γ be an (f,r)-piercing set of regions and P be a set of points. For any point $p \in P$, let $\Gamma(p)$ denote the subset of regions in Γ containing p. Then, the set system $(\Gamma, \{\Gamma(p) : p \in P\})$ (often called the* dual set system *defined by Γ and P) has shallow-cell complexity $O(r^2 f)$.*

Note that the above result implies the existence of *small* ϵ-nets for the dual set system and approximation algorithms with corresponding small approximation factors for the set cover problem on the *primal* set system $(P, \{\Gamma(p) : p \in P\})$. See the results in [8, 38].

▶ **Theorem 3.4.** *Let $(P, \mathcal{H})$ be a set system defined by a set P of points in the plane and an (f,r)-piercing set of regions $\mathcal{H}$, and let $w : \mathcal{H} \to \mathbb{N}$ be a weight function. Then,*
- **(i)** *the weighted Set Cover problem admits an $O(r^2 f)$-approximation algorithm.*
- **(ii)** *the weighted Discrete Independent Set problem admits an $\Omega(1/(r^2 f))$-approximation algorithm.*

4 (f,r)-Piercing regions

In this section, we describe properties of an (f,r)-piercing set of regions. Let Γ be an (f,r)-piercing set. We start with the notion of an *inscribed region* and use it to bound the number of components a region α is split into by regions $\beta \in \Gamma$.

▶ **Definition 4.1** (inscribed region). *We say that a region C is an inscribed region of another region α if $C \subseteq \alpha$ and ∂C touches $\partial\alpha$ at one or more (possibly infinite) points. The boundary ∂C consists of arcs that are along $\partial\alpha$ and arcs that are chords of α. We refer to the chords as the "sides" of C w.r.t. α and denote their number by $\textsc{sides}(C,\alpha)$.*

▶ **Lemma 4.2.** *Let α be any region with disjoint inscribed regions $C_1, C_2, \cdots, C_t$ where $t \geq 1$. Then, the number of connected components in $\alpha \setminus \bigcup_{j=1}^{t} C_j$ is $1 + \sum_{j=1}^{t} (\textsc{sides}(C_j,\alpha) - 1)$.*

Proof. We prove the lemma by induction on t. The base case is $t = 1$. Since each region of $\alpha \setminus C_1$ shares exactly one side of C_1, the number of connected components of $\alpha \setminus C_1$ is $\textsc{sides}(C_1,\alpha)$. Thus, the lemma holds for $t = 1$. Now, suppose that $t \geq 2$ and the lemma holds for smaller t. Since the regions $C_1, \cdots, C_t$ are disjoint, we can draw a chord σ of α which does not intersect any of the regions in $C_1, \cdots, C_t$ and splits α into two regions α_1 and α_2 so that one of the regions C_i lies in α_1 and remaining inscribed regions lie in α_2. By the induction hypothesis, the number of connected components in $\alpha_1 \setminus C_i$ is $m_1 = \text{SIDES}(C_i, \alpha)$ and the number of connected components in $\alpha_2 \setminus \bigcup_{j \neq i} C_j$ is $m_2 = 1 + \sum_{j \neq i} (\text{SIDES}(C_j, \alpha) - 1)$. This implies that the number of connected components of $\alpha \setminus \bigcup_{j=1}^{t} C_j$ is $m_1 + m_2 - 1$. We have a "-1" because the connected components of $\alpha_1 \setminus C_i$ and $\alpha_2 \setminus \bigcup_{j \neq i} C_j$ that have σ on their boundaries correspond to the same connected region in $\alpha \setminus \bigcup_{j=1}^{t} C_j$. The lemma follows. ◀

Let α and β be two regions whose boundaries intersect. Then, observe that every connected component C of $\alpha \cap \beta$ is an inscribed region of both α and β and furthermore $\text{SIDES}(C, \alpha) = \text{SIDES}(C, \beta)$.

▶ **Lemma 4.3.** *Let α and β be two regions whose boundaries intersect. Let $C_1, C_2, \cdots, C_t$ be the connected components in $\alpha \cap \beta$. Then,* $\text{FRAG}(\alpha, \beta) = 1 + \sum_{j=1}^{t} (\text{SIDES}(C_j, \alpha) - 1) = 1 + \sum_{j=1}^{t} (\text{SIDES}(C_j, \beta) - 1) = \text{FRAG}(\beta, \alpha)$.

Proof. Since $\text{SIDES}(C_j, \alpha) = \text{SIDES}(C_j, \beta)$ for each j, it suffices to prove that $\text{FRAG}(\alpha, \beta) = 1 + \sum_{j=1}^{t} (\text{SIDES}(C_j, \alpha) - 1)$ which follows directly from Lemma 4.2. ◀

The above lemma implies that $\text{FRAG}(\alpha, \beta) = \text{FRAG}(\beta, \alpha)$ for any α, β. If the boundaries of α and β don't intersect, by definition we have $\text{FRAG}(\alpha, \beta) = \text{FRAG}(\beta, \alpha) = 0$ and if they do then by Lemma 4.3, $\text{FRAG}(\alpha, \beta) = \text{FRAG}(\beta, \alpha)$.

▶ **Lemma 4.4.** *Let α be a region whose boundary intersects the boundaries of a set of disjoint regions $\beta_1, \beta_2, \cdots, \beta_t$. Then, the number of connected components in $X = \alpha \setminus \bigcup_{i=1}^{t} \beta_i$ is $1 + \sum_{i=1}^{t} (\text{FRAG}(\alpha, \beta_i) - 1)$. In particular, if at most r of the regions in $\{\beta_1, \cdots, \beta_t\}$ pierce α and $\text{FRAG}(\alpha, \beta_i) \leq f$ for each i, then the number of connected components in X is at most rf.*

Proof. Let $\mathcal{C}_i$ be set of connected components of $\alpha \cap \beta_i$ which are all inscribed regions of α and let $\mathcal{C} = \bigcup_{i=1}^{t} \mathcal{C}_i$. Since the β_i's are disjoint, the $\mathcal{C}_i$'s are disjoint sets. The number of connected components in X is the same as the number of connected components in $\alpha \setminus \bigcup_{C \in \mathcal{C}} C$ which by Lemma 4.2 is $1 + \sum_{C \in \mathcal{C}} (\text{SIDES}(C, \alpha) - 1)$ which is the same as $1 + \sum_{i=1}^{t} \sum_{C \in \mathcal{C}_i} (\text{SIDES}(C, \alpha) - 1) = 1 + \sum_{i=1}^{t} (\text{FRAG}(\alpha, \beta_i) - 1)$ by Lemma 4.3. Note that if β_i does not pierce α, then $\text{FRAG}(\alpha, \beta_i) - 1 = 0$ and therefore does not contribute to the sum above. Thus, if there are at most r regions piercing α, and $\text{FRAG}(\alpha, \beta_i) \leq f$ for all i, since the β_i are disjoint, it follows that the number of connected components of X is at most rf. ◀

We now show that applying lens bypassing does not increase the fragmentation number of the arrangement.

▶ **Lemma 4.5** (bypassing a minimal lens does not increase fragmentation number)**.** *Let Γ be a set of regions and let L be minimal lens (i.e., it does not contain another lens) in the arrangement. Let $\alpha, \beta \in \Gamma$ be the regions defining L. If we now bypass L by α i.e., modify α to $\alpha' = \alpha \setminus L_+$, then for any region $\gamma \in \Gamma \setminus \{\alpha\}$, $\text{FRAG}(\alpha', \gamma) \leq \text{FRAG}(\alpha, \gamma)$.*

Proof. We first show that $\text{FRAG}(\alpha', \beta) \leq \text{FRAG}(\alpha, \beta)$. To see this, note that the connected components of $\alpha' \cap \beta$ are the same as those in $\alpha \cap \beta$ except that the lens L we bypassed appears as a connected component in $\alpha \cap \beta$ but not in $\alpha' \cap \beta$. Since $\text{SIDES}(L, \alpha) = 1$, L does not contribute to the expression for $\text{FRAG}(\alpha, \beta)$ in Lemma 4.3. Thus if the boundaries of α' and β intersect, $\text{FRAG}(\alpha', \beta) = \text{FRAG}(\alpha, \beta)$. Even if $\partial\alpha'$ and $\partial\beta$ don't intersect, we have $0 = \text{FRAG}(\alpha', \beta) \leq \text{FRAG}(\alpha, \beta)$.

Now consider any region $\gamma \neq \beta$. The boundary of L consists of an arc σ_α of $\partial\alpha$ and an arc σ_β of $\partial\beta$. Since L is a minimal lens, the intersection of $\partial\gamma$ with L consists of disjoint arcs each with one endpoint on σ_α and the other endpoint on σ_β. Since bypassing L by α effectively moves σ_α to σ_β [4], the intersection of α with γ is combinatorially the same as the intersection of α' with γ. This implies that $\text{FRAG}(\alpha', \gamma) = \text{FRAG}(\alpha, \gamma)$. ◀

▶ **Definition 4.6** (Intersection graph of connected subgraphs and its depth). *Let H be any graph and let $\mathcal{C} = \{C_1, C_2, \cdots, C_k\}$ be a set of connected subgraphs of H. Then, the intersection graph defined by $\mathcal{C}$ is a graph G whose vertex set is $\mathcal{C}$ and there is an edge corresponding to any pair $\{C_i, C_j\}$ iff C_i and C_j share a vertex of H. We define the* depth *of the intersection graph as the maximum number of connected subgraphs in $\mathcal{C}$ containing any specific vertex in H.*

The Lemma below follows from Frederickson's decomposition [16].

▶ **Lemma 4.7.** *Let d and τ be fixed parameters. Let $\mathcal{G}$ be the family of graphs such that each graph $G \in \mathcal{G}$ is the depth d intersection graph of connected subgraphs in some planar graph H s.t. the number of vertices in H is at most τ times the number of vertices in G. Then, $\mathcal{G}$ admits a $(\lambda_1, \lambda_2, \lambda_3)$-decomposition where $\lambda_1 = O(d^{3/2}\tau)$, $\lambda_2 = O(d\tau)$ and $\lambda_3 = O(\sqrt{d})$.*

Proof. In order to show that $\mathcal{G}$ admits a $(\lambda_1, \lambda_2, \lambda_3)$-decomposition, we need to show that for any given parameter t and any graph $G = (V, E) \in \mathcal{G}$ with m vertices, we can obtain $X \subseteq V$ and a partition of $V \setminus X$ into $V_1, \cdots, V_\ell$ satisfying the conditions stated in Definition 2.10.

We obtain the required $X, V_1, \cdots V_\ell$ by applying the result of Frederickson [16] which shows that the family of planar graphs admits a (c_1, c_2, c_3)-decomposition for some constants c_1, c_2, c_3. Let $H = (U, F)$ be the planar graph with $n = |U| \leq \tau m$ s.t. G is the depth d intersection graph of connected subgraphs of H. For any $v \in V$, let C_v denote the connected subgraph in H corresponding to the vertex v in G. Applying the result of Frederickson [16] with parameter $t' = t/d$, we obtain sets $Y \subseteq U$ and a partition of $U \setminus Y$ into $U_1, \cdots, U_\ell$ s.t. $\ell \leq c_2 \frac{n}{t'}$, $|Y| \leq c_1 n/\sqrt{t'}$, $|U_i| \leq t'$ $\forall i$, $|N_H(U_i) \cap Y| \leq c_3\sqrt{t'}$ $\forall i$, and $N_H(U_i) \cap U_j = \emptyset$ $\forall i \neq j$.

Note that $\ell \leq c_2 d\tau \cdot m = \lambda_2 m$ with $\lambda_2 = c_2 d\tau$. We next define X as the set of all vertices in G whose corresponding connected subgraph in H contains one of the vertices in Y. Since any vertex in H is contained in at most d of the connected components, $|X| \leq d \cdot |Y| \leq c_1 \tau d^{3/2} \cdot \frac{m}{\sqrt{t}} = \lambda_1 \frac{m}{\sqrt{t}}$ with $\lambda_1 = c_1 d^{3/2}\tau$. For any i, we define V_i as the set of vertices in G whose corresponding connected subgraphs in H only contain vertices from U_i. Note that $V_1, \cdots, V_\ell$ is a partition of $V \setminus X$ since for any $v \in V$ note that if C_v does not contain any of the vertices in Y, then it can only contain vertices from exactly one group U_i (this relies on the fact that for $i \neq j$, $N(U_i) \cap U_j = \emptyset$). Since every vertex in U_i can be in at most d connected subgraphs, $|V_i| \leq d \cdot |U_i| \leq t$. Next, note that $|N_G(V_i) \cap X| \leq d \cdot |N_H(U_i) \cap Y| \leq c_3\sqrt{d} \cdot \sqrt{t} = \lambda_3\sqrt{t}$ with $\lambda_3 = c_3\sqrt{d}$. To see this, note that for any (v, x) s.t $v \in V_i$ and $x \in X$, C_x must contain a vertex of $N_H(U_i) \cap Y$ and each such vertex is contained in at most d connected subgraphs. Finally, $N_G(V_i) \cap V_j = \emptyset$ since for any $v \in V_i$ and $v' \in V_j$, C_v and $C_{v'}$ don't have a vertex in common in H – this follows from the fact that the vertices in C_v are from U_i and the vertices in $C_{v'}$ are from U_j and $U_i \cap U_j = \emptyset$. The lemma follows. ◀

We now define a two-colored version of an (f, r)-piercing regions with additional properties. This will be useful in the analysis of the algorithms in Section 5.

[4] more precisely a curve arbitrarily close to it

▶ **Definition 4.8** (special (f,r)-piercing regions). *A* special *(f,r)-piercing set of regions is an (f,r)-piercing set where each region is colored either red or blue, and (i) any two regions of the same color are either disjoint or piercing, (ii) there are no* monochromatic *lenses i.e., lenses (not necessarily minimal) formed by pairs of regions of the same color and (iii) every* bi-chromatic *lens (i.e., a lens formed by a pair of regions of different color) is a minimal lens.*

For any special (f,r)-piercing set Γ, let $I(\Gamma)$ denote the *bipartite intersection graph* (Γ, E) where there is an edge between any intersecting pair of regions of different colors in Γ.

▶ **Lemma 4.9.** *Let Γ be an (f,r)-piercing set of regions in which each region is colored either red or blue and let P be a set of points s.t. Γ and P are in general position. Then, we can construct a special (f,r)-piercing set Γ' in which corresponding to every $\gamma \in \Gamma$, there is a region $\gamma' \subseteq \gamma$ of the same color as γ so that the following condition holds for any $p \in P$. If p is contained in at least one red and one blue region in Γ, it is also contained in at least one red and one blue region in Γ'.*

Proof. We will do a sequence of minimal lens bypassing operations so that for any point $p \in P$ the following *coverage condition* holds: if p is contained in at least one red and one blue region before the bypassing, it is still contained in at least one red and one blue region after the bypassing. After such simplifications, we will show that the resulting arrangement is a special (f,r)-piercing set. Note that in any lens bypassing any region γ is either not modified or modified to a region $\gamma' \subseteq \gamma$. Also, note that by Lemma 4.5, the set of regions remains (f,r)-piercing at all times. First, note that upon bypassing any minimal lens formed by two regions of the same color, regardless of the region modified, the coverage condition holds. Therefore we can bypass all such *monochromatic* minimal lenses. Now consider a *bi-chromatic* minimal lens formed by a red and a blue region which is contained in a third region (red or blue). We can also bypass such lenses while maintaining the coverage condition by making sure that the region modified during the bypassing is the one whose color matches the third region. After we have bypassed all minimal lenses of the above two types, observe that we cannot have any monochromatic lenses (even non-minimal). To see this, consider a minimal monochromatic lens i.e., a monochromatic lens which does not contain another monochromatic lens. Such a lens is either a minimal lens or it contains a bi-chromatic lens. In the former case, the lens should have been bypassed, and in the latter case, the bi-chromatic lens is contained in a third region and should have been bypassed. This implies that any regions of the same color are either disjoint or are piercing. The lemma follows. ◀

▶ **Definition 4.10** (simplified special (f,r)-piercing set). *A simplified special (f,r)-piercing set of regions is a special (f,r)-piercing set s.t. (i) any two regions in Γ that pierce each other don't form any lenses, and (ii) any two regions in Γ that don't pierce intersect in at most one lens.*

▶ **Lemma 4.11** (obtaining a simplified special (f,r)-piercing set). *Given any special (f,r)-piercing set Γ, we can obtain a simplified set Γ' in which for each region $\gamma \in \Gamma$, there is a region $\gamma' \subseteq \gamma$ of the same color as γ and the following condition holds. For any two regions $\gamma_1, \gamma_2 \in \Gamma$ of different colors intersect iff γ'_1 and γ'_2 intersect.*

Proof. Since Γ is a special (f,r)-piercing set, the arrangement of regions in Γ does not contain any monochromatic lenses and all bi-chromatic lenses are minimal. We repeatedly bypass bi-chromatic minimal lenses formed by regions that also intersect outside the lens as this does not violate the conditions of the lemma. Thus if any lens is left it is the only cell in the intersection of two regions forming it. In particular, the two regions cannot be piercing and they cannot be intersecting in any other lens. ◀

▶ **Lemma 4.12.** *Let Γ be a simplified special (f,r)-set of regions. The number of cells in the arrangement of the regions in Γ is $O(r^2 fn)$ where $n = |\Gamma|$.*

Proof. It suffices to show that the number of intersections of the boundaries of pairs of regions in Γ is $O(r^2 fn)$. In order to do this, we partition the red regions into $r+1$ sets $\mathcal{R}_1, \cdots, \mathcal{R}_{r+1}$ s.t. the regions in each set are disjoint. This is because regions of the same color are either piercing or disjoint and no region is pierced by more than r other regions. Similarly, we partition the blue regions into $r+1$ sets $\mathcal{B}_1, \cdots, \mathcal{B}_{r+1}$ s.t. the regions in each set are disjoint. Overall, we have $2(r+1)$ sets of disjoint regions. Consider any two of these families A and B (which may or may not be of the same color).

We next argue that the number of intersections among the regions in $A \cup B$ is $O(rfn)$ where $m = |A| + |B|$. To see this, imagine replacing each region $\alpha \in A$ by the connected components in $\alpha \setminus \bigcup_{\beta \in B} \beta$. The number of components any α is split into is at most rf by Lemma 4.4. The overall number of regions now is at most $m' = |A| \cdot rf + |B| \leq rfm$ which are interior disjoint but certain pairs touch along their boundaries.

Consider the *touching graph* defined by these regions in which there is a vertex corresponding to each region and two vertices are adjacent iff the boundaries of their corresponding regions overlap in an arc of non-zero length. Note that this graph is planar and hence, the number of edges in this graph is at most $3m'$. Here it is important to note also that the touching graph does not have multi-edges due to the assumption that two piercing regions do not form any lenses and two non-piercing regions form at most one lens. Each edge in the touching graph corresponds to two intersection points among boundaries of regions in $A \cup B$, showing that the number of intersections among the regions in $A \cup B$ is at most $3rfm$.

Since each set $A \in \mathcal{S} = \{\mathcal{R}_1, \cdots, R_{r+1}\} \cup \{\mathcal{B}_1, \cdots, B_{r+1}\}$ appears in $O(r)$ pairs of the form $(A,B) \in \binom{\mathcal{S}}{2}$ and $\sum_{i=1}^{r+1} |\mathcal{R}_i| + |\mathcal{B}_i| = n$, the overall number of intersections among the boundaries of regions in Γ is $O(r^2 fn)$ which in turn implies that the number of cells in the arrangement of regions in Γ is also $O(r^2 fn)$. ◀

▶ **Lemma 4.13.** *For any fixed constants f, r, the family of bipartite intersection graphs $\mathcal{I} = \{I(\Gamma) : \Gamma$ is a special (f,r)-piercing set$\}$ admits a $(\lambda_1, \lambda_2, \lambda_3)$-decomposition where $\lambda_1 = O(r^{3.5} f), \lambda_2 = O(r^3 f)$, and $\lambda_3 = O(\sqrt{r}))$.*

Proof. Consider a special (f,r)-set Γ. By Lemma 4.11, we can assume without loss of generality, that Γ is a simplified special (f,r)-set. By Lemma 4.12, the number of cells in the arrangement of the regions in Γ is $O(r^2 fn)$ Now consider the dual arrangement graph (Definition 2.4) H of the regions in Γ. Note that H is planar and the cells contained in any region $\gamma \in \Gamma$ correspond to a connected subgraph of H. Note also that any cell is contained in at most $r+1$ red regions and at most $r+1$ blue regions and therefore is in at most $2(r+1)$ regions overall. In other words, each graph $G \in \mathcal{I}$ is the depth $d = 2(r+1)$ intersection graph of connected subgraphs in a planar graph H having at most τn vertices where $\tau = O(r^2 f)$. By Lemma 4.7, $\mathcal{I}$ admits a $(\lambda_1, \lambda_2, \lambda_3)$-decomposition where $\lambda_1 = O(d^{3/2}\tau) = O(r^{3.5} f), \lambda_2 = O(d\tau) = O(r^3 f)$, and $\lambda_3 = O(\sqrt{d}) = O(\sqrt{r})$. ◀

5 Discrete Independent Set and Set Cover

Let Γ be a set of (f,r)-piercing regions and let P be a set of points in the plane.

In this section, we define the local search algorithms and show that they yield a PTAS for the Discrete Independent Set and Set Cover problems.

t-Local Search

For any given parameter t, the algorithm is the following. We always maintain a feasible solution S which in the beginning is chosen arbitrarily. At any stage, we check if it is possible to replace some subset of the regions T by another subset of regions T' so that $|T|, |T'| \leq t$ and $S' = (S \setminus T) \cup T'$ is a feasible solution with a better objective value than S. If this is possible, we replace S by S' and continue. If such a *local improvement* is not possible, we say that the current solution is *t-locally optimal* in which case, we simply return the current solution S. A straightforward implementation of the t-local search algorithm has running time $O(n^{O(t)})$ where n is the total number of regions.

▶ **Theorem 3.1** (Discrete Independent Set). *The t-local search algorithm with $t = Cr^7 f^2/\epsilon^2$ for an appropriately large C and $\epsilon \in (0, 1]$ yields a $(1-\epsilon)$-approximation for the Discrete Independent Set problem defined by a set of (f, r)-piercing regions Γ and a set of points P. The algorithm runs in time $O(n^{O(t)})$.*

Proof. Let Opt be an optimal solution and Local be the solution returned by the t-local search algorithm. For the purpose of the analysis, it helps to assume that Opt and Local are disjoint. This can be done by removing the regions common to both and then comparing their sizes. Formally, let $I = \text{Opt} \cap \text{Local}$ and let $R = \text{Opt} \setminus I$ and $B = \text{Local} \setminus I$. Let $P' \subseteq P$ be the set of points not covered by the regions in I. Note that R and B are disjoint and are the optimal and t-locally optimal solutions respectively with respect to P'. For the following arguments, we will color the regions in R red and the regions in B blue.

By applying Lemma 4.9, to the set of regions $\Gamma = R \cup B$ and the set of points P', we obtain a special (f, r)-set Γ' where for each region $\gamma \in \Gamma$, there is a corresponding *shrunk* region $\gamma' \subseteq \gamma$ in Γ'. Let $G = (R \cup B, E)$ be a bipartite graph in which there is an edge (r, b) for some $r \in R$ and $b \in B$ iff r' and b' (the *shrunk* regions corresponding to r and b) intersect. Note that G is isomorphic to the bipartite intersection graph of the regions in Γ'.

Note that every point $p \in P'$ is covered by at most one red region in R and one blue region in B and furthermore if p is contained in a red region r as well as a blue region b, then there an edge (r, b) in G. Thus, the following *local search condition* holds: $(B \cup \tilde{R}) \setminus N_G(\tilde{R})$ is a feasible solution for any $\tilde{R} \subseteq R$ where $N_G(\tilde{R})$ denotes the set of blue neighbors of the regions in $\tilde{R}$ in G.

By Lemma 4.13, G has a $(\lambda_1, \lambda_2, \lambda_3)$-decomposition where $\lambda_1 = O(r^{3.5} f), \lambda_2 = O(r^3 f)$, and $\lambda_3 = O(\sqrt{r})$ i.e., there is a set $X \subseteq \Gamma'$ with $|X| \leq \lambda_1 m/\sqrt{t}$ and a partition of $V \setminus X$ into $\ell \leq \lambda_2 m/t$ sets, where $m = |\Gamma|$, s.t. (i) $|V_i| \leq t$ for all i, (ii) $|N_G(V_i) \cap X| \leq \lambda_3 \sqrt{t}$ for all i, and (iii) $N_G(V_i) \cap V_j = \emptyset$ for $i \neq j$. For $i = 1, \cdots, \ell$, let $B_i = V_i \cap B$ and $R_i = V_i \cap R$. We now show that for any i, $|B_i| \geq |R_i| - \lambda_3 \sqrt{t}$. For contradiction suppose that we have $|B_i| < |R_i| - \lambda_3 \sqrt{t}$ for some i. By the local search condition: $S = (B \cup R_i) \setminus N_G(R_i)$ is a feasible solution. Since $|N_G(R_i)| \leq |B_i| + \lambda_3 \sqrt{t} < |R_i|$, S is an improved feasible solution obtained from B by adding $|R_i| \leq t$ elements and removing $|N_G(R_i)| < |R_i|$ elements contradicting the assumption that B is t-locally optimal. Now, $|B| \geq \sum_i (|R_i| - \lambda_3 \sqrt{t}) \geq \sum_i |R_i| - \lambda_2 \frac{m}{t} \cdot \lambda_3 \sqrt{t} \geq |R| - |X| - \lambda_2 \lambda_3 \frac{m}{\sqrt{t}} \geq |R| - (\lambda_1 + \lambda_2 \lambda_3) \frac{m}{\sqrt{t}}$ where $m = |R| + |B|$ from which it follows that $|B| \geq (1 - \epsilon)|R|$ for sufficiently large C. This implies that $|\text{Local}| \geq (1 - \epsilon)|\text{Opt}|$. The theorem follows. ◀

The proof of the following theorem is very similar to the proof of Theorem 3.1.

▶ **Theorem 3.2** (Set Cover). *The t-local search algorithm with $t = Cr^7 f^2/\epsilon^2$ for an appropriately large C and $\epsilon \in (0, 1]$ yields a $(1+\epsilon)$-approximation for the Set Cover problem defined by a set of (f, r)-piercing regions Γ and a set of points P. The algorithm runs in time $n^{O(t)}$.*

Proof. The proof is very similar to the proof of Theorem 3.1 and we use the notation defined there. Since every point $p \in P'$ is covered by at least one red region and one blue region, the following *local search condition* holds: $(B \setminus B') \cup N_G(B')$ is a feasible solution for any $B' \subseteq B$. As in the proof of Theorem 3.1, we apply Lemma 4.13 and obtain the sets X, B_i and R_i. We next show that for any i, $|B_i| \le |R_i| + \lambda_3\sqrt{t}$. For contradiction, suppose that this is not the case for some i. By the local search condition, $S = (B \setminus B_i) \cup N_G(B_i)$ is a feasible solution. Since $|N_G(B_i)| \le |R_i| + \lambda_3\sqrt{t} < |B_i|$, S is an improved feasible solution obtained by removing $|B_i| \le t$ elements from B and adding $|N_G(B_i)| < |B_i|$ elements, contradicting the assumption that B is t-locally optimal. Now $|B| \le |X| + \sum_i |B_i| \le |X| + \sum_i (|R_i| + \lambda_3\sqrt{t}) \le \lambda_1 \frac{m}{\sqrt{t}} + |R| + \lambda_2\lambda_3 \frac{m}{\sqrt{t}}$ where $m = |R| + |B|$. Assuming C is sufficiently large, we obtain $|B| \le (1+\epsilon)|R|$ which implies that $|\textsc{Local}| \le (1+\epsilon)|\textsc{Opt}|$. The theorem follows. ◀

Now, we show that the *dual set system* defined by a set of points and an (f, r)-piercing set of regions has bounded shallow-cell complexity.

▶ **Theorem 3.3.** *Let Γ be an (f, r)-piercing set of regions and P be a set of points. For any point $p \in P$, let $\Gamma(p)$ denote the subset of regions in Γ containing p. Then, the set system $(\Gamma, \{\Gamma(p) : p \in P\})$ (often called the* dual set system *defined by Γ and P) has shallow-cell complexity $O(r^2 f)$.*

Proof. We first prove that if Γ is a special (f, r)-set then the number of combinatorially distinct *bi-chromatic* cells of depth 2 (i.e., cells contained in exactly one red and one blue region) in the arrangement Γ is $O(r^2 fn)$. We then apply a standard argument due Clarkson and Shor [10] to prove the lemma. Let us first assume that Γ is a special (f, r)-set. Since we are only interested in combinatorially distinct bi-chromatic cells, we may assume by Lemma 4.11 that Γ is a simplified special (f, r)-set and by Lemma 4.12 the number of such cells is $O(r^2 fn)$. Now, assume that Γ is any (f, r)-piercing set. Our goal is to bound the number of combinatorially distinct cells of depth $\le k$ in the arrangement Γ. Since every cells has a vertex, it suffices to count the vertices of depth $\le k$ contained in a combinatorially distinct set of regions. Let $N_{\le k}$ denote this number. Let $\Gamma' \subseteq \Gamma$ be set obtained by picking each region in Γ independently with probability $p = 1/k$ and coloring it red or blue uniformly at random. The probability that a vertex of depth $\le k$ appears as a bi-chromatic vertex (a vertex formed by intersection of the boundaries of a red and a blue region but not contained in any other region) in the arrangement Γ' is at least $\frac{1}{2}p^2(1-p)^k \ge \frac{1}{2}p^2/e = 1/(2ek^2)$. Hence, in expectation, we get at least $N_{\le k}/(2ek^2)$ such vertices. By Lemma 4.9, the set Γ' can be modified to a special (f, r)-piercing set with the same set of combinatorially distinct bi-chromatic cells. This means that number of bi-chromatic vertices in the arrangement Γ' is at most $O(r^2 f|\Gamma'|)$ which is $O(r^2 fn/k)$ in expectation. Thus, $N_{\le k}/(2ek^2) \le r^2 fn/k$ implying that $N_{\le k}$ is $O(kr^2 fn)$. ◀

6 Conclusion

We have introduced piercing complexity as a natural topological parameter for geometric optimization and shown that bounded piercing enables PTASs for fundamental packing and covering problems. Specifically, for (f, r)-piercing regions, we obtained PTASs for unweighted Set Cover and Discrete Independent Set problems, along with $O(r^2 f)$-approximations for their weighted variants. These results generalize known algorithms for non-piercing regions and yield improved guarantees for important special cases, including axis-parallel rectangles and line segments under bounded piercing.

Our results build on the local-search framework developed for non-piercing regions [5, 35] but extend it to a significantly broader class of instances. Several important questions remain open. First, does a PTAS exist for the Hitting Set problem with (f, r)-piercing regions? This appears non-trivial and likely requires deeper structural insights. Second, while known hardness constructions show that unbounded piercing can be an obstruction to PTASs, it remains unclear whether the fragmentation number f is required to be bounded, or whether a weaker condition would suffice.

More broadly, piercing complexity captures a fundamental source of hardness that existing parameters − VC-dimension, shallow-cell complexity, fatness, and density − fail to measure. These parameters control combinatorial richness or geometric overlap but do not capture how intersections topologically fragment regions. By quantifying this fragmentation, piercing complexity directly measures a structural property that obstructs local-search techniques. In this sense, (f, r)-piercing gives a broad structural extension of the non-piercing regime under which PTASs remain attainable, providing a unifying framework that explains both classical hardness constructions and previously isolated positive results. Finally, a finer structural understanding of piercing interactions may prove useful in resolving longstanding open problems, such as the Set Cover and Independent Set problems for axis-parallel rectangles, and the problems with line segments and points in the plane.

References

1 Anna Adamaszek and Andreas Wiese. A QPTAS for maximum weight independent set of polygons with polylogarithmically many vertices. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 645–656, 2014. doi:10.1137/1.9781611973402.49.

2 Pankaj K. Agarwal, Marc van Kreveld, and Subhash Suri. Label placement by maximum independent set in rectangles. *Computational Geometry*, 11(3):209–218, 1998. doi:10.1016/S0925-7721(98)00028-5.

3 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*, volume 7748 of *Lecture Notes in Computer Science*, pages 89–100. Springer, 2013. doi:10.1007/978-3-642-36065-7_10.

4 Jozsef Balogh and Jozsef Solymosi. On the number of points in general position in the plane. *Discrete Analysis*, pages 1–20, 2018. doi:10.19086/da.4438.

5 Aniket Basu Roy, Sathish Govindarajan, Rajiv Raman, and Saurabh Ray. Packing and covering with non-piercing regions. *Discrete & Computational Geometry*, 2018.

6 Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995. doi:10.1007/BF02570718.

7 Timothy M. Chan and Elyot Grant. Exact algorithms and APX-hardness results for geometric packing and covering problems. *Computational Geometry Theory and Applications*, 47(2):112–124, 2014. doi:10.1016/j.comgeo.2012.04.001.

8 Timothy M. Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, pages 1576–1585, 2012. doi:10.1137/1.9781611973099.125.

9 Timothy M. Chan and Sariel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete and Computational Geometry*, 48(2):373–392, 2012. doi:10.1007/S00454-012-9417-5.

10 Kenneth L. Clarkson and Peter W. Shor. Applications of random sampling in computational geometry, II. *Discrete & Computational Geometry*, 4(5):387–421, 1989. doi:10.1007/BF02187740.

11 Kenneth L. Clarkson and Kasturi R. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007. doi:10.1007/s00454-006-1273-8.

12 Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized algorithms*, volume 5(4). Springer, 2015. doi:10.1007/978-3-319-21275-3.

13 Thomas Erlebach and Erik Jan van Leeuwen. PTAS for weighted set cover on unit squares. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 13th International Workshop, APPROX 2010, and 14th International Workshop, RANDOM 2010, Barcelona, Spain, September 1-3, 2010. Proceedings*, pages 166–177, 2010. doi:10.1007/978-3-642-15369-3_13.

14 Guy Even, Dror Rawitz, and Shimon Shahar. Hitting sets when the vc-dimension is small. *Inf. Process. Lett.*, 95(2):358–362, 2005. doi:10.1016/j.ipl.2005.03.010.

15 Uriel Feige. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998. doi:10.1145/285055.285059.

16 Greg N Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM Journal on Computing*, 16(6):1004–1022, 1987. doi:10.1137/0216064.

17 Waldo Gálvez, Arindam Khan, Mathieu Mari, Tobias Mömke, Madhusudhan Reddy Pittu, and Andreas Wiese. A 3-approximation algorithm for maximum independent set of rectangles. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 894–905. SIAM, 2022. doi:10.1137/1.9781611977073.38.

18 Matt Gibson and Imran A. Pirwani. Algorithms for dominating set in disk graphs: Breaking the logn barrier. In Mark de Berg and Ulrich Meyer, editors, *Algorithms – ESA 2010*, pages 243–254, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.

19 Daniel Golovin, Anupam Gupta, Amit Kumar, and Kanat Tangwongsan. All-norms and all-l_p-norms approximation algorithms. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (2008)*, pages 199–210. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2008. doi:10.4230/LIPIcs.FSTTCS.2008.1753.

20 Sariel Har-Peled. *Geometric approximation algorithms*. Number 173 in Mathematical Surveys and Monographs. American Mathematical Soc., 2011.

21 Sariel Har-Peled and Kent Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. In *Algorithms – ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14–16, 2015, Proceedings*, pages 717–728, 2015. doi:10.1007/978-3-662-48350-3_60.

22 Sariel Har-Peled and Kent Quanrud. Approximation algorithms for polynomial-expansion and low-density graphs. *SIAM J. Comput.*, 46(6):1712–1744, 2017. doi:10.1137/16M1079336.

23 David Haussler and Emo Welzl. epsilon-nets and simplex range queries. *Discrete & Computational Geometry*, 2:127–151, 1987. doi:10.1007/BF02187876.

24 Dorit S. Hochbaum and Wolfgang Maass. Approximation schemes for covering and packing problems in image processing and vlsi. *J. ACM*, 32(1):130–136, 1985. doi:10.1145/2455.214106.

25 J Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182(1):105–142, 1999.

26 Bruno Jartoux and Nabil H. Mustafa. Optimality of Geometric Local Search. In *34th International Symposium on Computational Geometry (SoCG 2018)*, Leibniz International Proceedings in Informatics (LIPIcs), pages 48:1–48:15, 2018. doi:10.4230/LIPIcs.SOCG.2018.48.

27 Katarzyna Kowalska and Michał Pilipczuk. Parameterized and approximation algorithms for coverings points with segments in the plane. In *41st International Symposium on Theoretical Aspects of Computer Science*, 2024.

28 Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, Jie Xue, and Meirav Zehavi. Subexponential parameterized algorithms on disk graphs. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2005–2031. SIAM, 2022.

29 Laszlo Lovász. On the ratio of optimal integral and fractional covers. *Discrete Math*, 13:383–390, 1975. doi:10.1016/0012-365X(75)90058-8.

30 Jiří Matoušek. *Lectures on Discrete Geometry.* Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2002.

31 Joseph S. B. Mitchell. Approximating maximum independent set for rectangles in the plane. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 339–350. IEEE, 2021. doi:10.1109/FOCS52979.2021.00042.

32 Nabil H. Mustafa. *Sampling in combinatorial and geometric set systems*, volume 265. American Mathematical Society, 2022.

33 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010. doi:10.1007/S00454-010-9285-9.

34 János Pach and Gábor Tardos. Tight lower bounds for the size of epsilon-nets. In *Proceedings of the twenty-seventh annual symposium on Computational geometry*, pages 458–463, 2011. doi:10.1145/1998196.1998271.

35 Rajiv Raman and Saurabh Ray. Constructing planar support for non-piercing regions. *Discrete and Computational Geometry*, 64(3):1098–1122, 2020. doi:10.1007/S00454-020-00216-W.

36 Rajiv Raman and Saurabh Ray. On the geometric set multicover problem. *Discret. Comput. Geom.*, 68(2):566–591, 2022. doi:10.1007/s00454-022-00402-y.

37 Rajiv Raman and Karamjeet Singh. Supports for outerplanar and bounded treewidth graphs, 2025. doi:10.48550/arXiv.2504.05039.

38 Kasturi Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In *Proceedings of the 42nd ACM Symposium on Theory of Computing*, pages 641–648, 2010. doi:10.1145/1806689.1806777.

39 Bang Wang. Coverage problems in sensor networks: A survey. *ACM Computing Surveys (CSUR)*, 43(4):1–53, 2011. doi:10.1145/1978802.1978811.