Approximation Algorithms for Maximum Independent Set of Pseudo-Disks*

Timothy M. Chan† Sariel Har-Peled‡

December 3, 2008

Abstract

We present approximation algorithms for maximum independent set of pseudo-disks in the plane, both in the weighted and unweighted cases. For the unweighted case, we prove that a local search algorithm yields a PTAS. For the weighted case, we suggest a novel rounding scheme based on an LP relaxation of the problem, that leads to a constant-factor approximation.

Most previous algorithms for maximum independent set (in geometric settings) relied on packing arguments that are not applicable in this case. As such, the analysis of both algorithms requires some new combinatorial ideas, which we believe to be of independent interest.

---

*Alternative titles for this paper include: "On the pseudo-disks march to maximum independence" and "Lighter reading for the second coming of the great depression."

†School of Computer Science; University of Waterloo; 200 University Ave West; Waterloo, Ontario N2L 3G1; Canada; tmchan@uwaterloo.ca; http://www.cs.uwaterloo.ca/~tmchan/.

‡Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@uiuc.edu; http://www.uiuc.edu/~sariel/.

1 Introduction

Let $\mathsf{F} = \{f_1, \ldots, f_n\}$ be a set of n objects in the plane, with weights $w_1, w_2, \ldots, w_n > 0$, respectively. In this paper, we are interested in the problem of finding an independent set of maximum weight. Here a set of objects is ***independent***, if no pair of objects intersect.

A natural approach to this problem, is to build an ***intersection graph*** $G = (V, E)$, where the objects form the vertices, and two objects are connected by an edge if they intersect, and weights are associated with the vertices. We want the maximum independent set in G. This is of course an NP-COMPLETE problem, and it is known that no approximation factor is possible within $|V|^{1-\varepsilon}$ for any $\varepsilon > 0$ if $\mathsf{NP} \neq \mathsf{ZPP}$ [Has96]. In fact, even if the maximum degree of the graph is bounded by 3, no PTAS is possible in this case [BF99].

In geometric settings, better results are known. If the objects are fat (e.g., disks and squares), PTASs are known. One approach [Cha03, EJS05] relies on a hierarchical spatial subdivision, such as a quadtree, combined with dynamic programming techniques [Aro98]; it works even in the weighted case. Another approach [Cha03] relies on a recursive application of a nontrivial generalization of the planar separator theorem [LT79, SW98]; this approach is limited to the unweighted case. If the objects are not fat, only weaker results are known. For the problem of finding a maximum independent set of unweighted axis-parallel rectangles, an $O(\log \log n)$-approximation algorithm was very recently given by Chalermsook and Chuzhoy [CC09]. For line segments, a roughly $O(\sqrt{\mathrm{Opt}})$-approximation is known [AM06].

In this paper we are interested in the problem of finding a large independent set in a set of weighted or unweighted pseudo-disks. A set of objects is a collection of ***pseudo-disks***, if the boundary of every pair of them intersects at most twice. This case is especially intriguing because previous techniques seem powerless: it is unclear how one can adapt the quadtree approach [Cha03, EJS05] or the generalized separator approach [Cha03] for pseudo-disks.

Even a constant-factor approximation in the unweighted case is not easy. Consider the most obvious greedy strategy for disks (or fat objects): select the object $f_i \in \mathsf{F}$ of the smallest radius, remove all objects that intersect f_i from F, and repeat. This is already sufficient to yield a constant-factor approximation by a simple packing argument [EKNS00, MBH+95]. However, even this simplest algorithm breaks down for pseudo-disks—as pseudo-disks are defined "topologically", how would one define the "smallest" pseudo-disk in a collection?

An independent set via local search. Nevertheless, we are able to prove that a different strategy can yield a constant-factor approximation for unweighted pseudo-disks: local search. In the general settings, local search was used to get (roughly) a $\Delta/4$ approximation to independent set, where Δ is the maximum degree in the graph, see [Hal98] for a survey. In the geometric settings, Agarwal and Mustafa [AM06, Lemma 4.2] had a proof that a local search algorithm gives a constant-factor approximation for the special case of pseudo-disks that are rectangles; their proof does not immediately work for arbitrary pseudo-disks. Our proof provides a generalization of their lemma.

In fact, we are able to do more: we show that local search can actually yield a PTAS for unweighted pseudo-disks! This gives us by accident a new PTAS for the special case of disks and squares. Though the local-search algorithm is slower than the quadtree-based PTAS in these special cases [Cha03], it has the advantage that it only requires the intersection graph as input, not its geometric realization; previously, an algorithm with this property was only known in further special cases, such as unit disks [NHK05]. Our result uses the planar separator theorem, but in contrast to the separator-based method in [Cha03], a standard version of the theorem suffices and is needed only in the analysis, not in the algorithm itself.

Planar graphs are special cases of disk intersection graphs, and so our result applies. Of course, PTASs for planar graphs have been around for quite some time [LT79, Bak94], but the fact that a simple local search algorithm already yields a PTAS for planar graphs is apparently not well known, if it was known at all.

This strategy, unfortunately, works only in the unweighted case.

An independent set via LP. It is easy to extract a large independent set from a sparse unweighted graph. For example, greedily, we can order the vertices from lowest to highest degree, and pick them one by one into the independent set, if none of its neighbors was already picked into the independent set. Let d_G be the average degree in G. Then a fraction of the vertices have degree $O(d_G)$, and the selection of such a vertex can eliminate $O(d_G)$ candidates. Thus, this yields an independent set of size $\Omega(d_G)$. Alternatively, for better constants, we can order the vertices by a random permutation and do the same. Clearly, the probability of a vertex v to be included in the independent set is $1/(d(v)+1)$. An easy calculation leads to Turán's theorem, which states that any graph G has an independent set of size $\geq n/(d_G+1)$ [AS00].

Now, our intersection graph G may not be sparse. We would like to "sparsify" it, so that the new intersection graph is sparse and the number of vertices is close to the size of the optimal solution. Interestingly, we show that this can be done by solving the LP relaxation of the independent set problem. The relaxation provides us with a fractional solution, where every object f_i has value $x_i \in [0,1]$ associated with it. Rounding this fractional solution into a feasible solution is not a trivial task, as no such scheme exists in the general case. To this end, we prove a technical lemma (see Lemma 4.1) that shows that the total sum of terms of the form $x_i x_j$, over pairs $f_i f_j$ that intersect, is bounded by the boundary complexity of the union of $\mathcal{E}$ objects of F, where $\mathcal{E}$ is the size of the fractional solution. The proof contains a nice application of the standard Clarkson technique [CS89], applied in a way that is somewhat different from usual.

This lemma implies that on average, if we pick f_i into our random set of objects, with probability x_i, then the resulting intersection graph would be sparse. This is by itself sufficient to get a constant-factor approximation for the unweighted case. For the weighted case, we order the objects in decreasing orders by their weight. We argue that by doing the selection into the independent set according to this ordering, if an object has many objects intersecting it, then because of the above lemma, we can charge it to heavier objects that were already selected. This leads to a constant-factor approximation for weighted pseudo-disks.

Linear union complexity. Our LP analysis works more generally for any class of objects with linear union complexity. We assume that the boundary of the union of any k of these objects has at most ϱk vertices, for some fixed ϱ. For pseudo-disks, the boundary of the union is made out of at most $6n-12$ arcs, implying $\varrho = 6$ in this case [KLPS86].

A family F of simply connected regions bounded by simple closed curves in general position in the plane is ***k-admissible*** (with k even) if for any pair $f_i, f_j \in \mathsf{F}$, we have: (i) $f_i \setminus f_j$ and $f_j \setminus f_i$ are connected, and (ii) their boundary intersect at most k times. Whitesides and Zhao [WZ90] showed that the union of such n objects has at most $3kn - 6$ arcs; that is, $\varrho = 3k$. So, our LP analysis applies to this class of objects as well. For more results on union complexity, see the sermon by Agarwal *et al.* [APS08].

Our local-search PTAS works more generally for unweighted admissible regions. For an arbitrary class of unweighted objects, local search still yields a constant-factor approximation.

Discussion. Local search and LP relaxation are of course staples in the design of approximation algorithms, but are not seen as often in computational geometry. Our main contribution lies in the fusion of these approaches with combinatorial geometric techniques.

LP relaxation has been used before, notably, in Chalermsook and Chuzhoy's recent breakthrough in the case of axis-parallel rectangles [CC09], but their analysis is quite complicated. Although rectangles do not have linear union complexity in general, we observe in Appendix A that a variant of our approach can yield a readily accessible proof of a sublogarithmic $O(\log n / \log \log n)$ approximation factor for rectangles, even in the weighted case, where previously only logarithmic approximation is known [AvKS98, BDMR01, Cha04] (Chalermsook and Chuzhoy's result is better but currently is applicable only to unweighted rectangles).

In a sense, one can view our results as complementary to the known results on approximate geometric set cover by Brönnimann and Goodrich [BG95] and Clarkson and Varadarajan [CV07]. They consider the problem of finding the minimum number of objects in F to cover a given point set. Their results imply a constant-factor approximation for families of objects with linear union complexity, for instance. One version of their approaches is indeed based on LP relaxation [Lon01, ERS05]. The "dual" hitting set problem is to find the minimum number of points to pierce a given set of objects. Brönnimann and Goodrich's result combined with a recent result of Pyrga and Ray [PR08] also implies a constant-factor approximation for pseudo-disks for this piercing problem. The piercing problem is actually the dual of the independent set problem (this time, we are referring to linear programming duality). We remark, however, that the rounding schemes for set cover and piercing are based on different combinatorial techniques, namely, ε-nets, which are not sufficient to deal with independent set (one obvious difference is that independent set is a maximization problem).

In Theorem 4.6, we point out a combinatorial consequence of our LP analysis: for any collection of unweighted pseudo-disks, the ratio of the size of the minimum piercing set to the size of maximum independent set is at most a constant. (It is easy to see that the ratio is always at least 1; for disks or fat objects, it is not difficult to obtain a constant upper bound by packing arguments.) This result is of independent interest; for example, getting tight bounds on the ratio for axis-parallel rectangles is a long-standing open problem.

2 Preliminaries

In the following, we have a set of n regions in the plane F, such that the union complexity of any subset $X \subseteq \mathsf{F}$ is bounded by $\varrho\,|X|$, where ϱ is a constant. Here, the ***union complexity*** of X is the number of arcs on the boundary of the union of the objects of X. Let $\mathcal{A}(\mathsf{F})$ denote the arrangement of F, and $\mathcal{V}(\mathsf{F})$ denote the set of vertices of $\mathcal{A}(\mathsf{F})$.

In the following, we assume that deciding if two regions intersects takes a constant time.

3 Approximation using local search – unweighted case

3.1 The algorithm

In the unweighted case, we may assume that no object is fully contained in another.

We say that a subset L of F is ***b*-locally optimal** if T is an independent set and one cannot obtain a larger independent set from T by deleting at most b objects and inserting at most $b+1$ objects of F.

Our algorithm for the unweighted case simply returns a b-locally optimal solution for a suitable constant b, by performing a local search. We start with $\mathsf{L} \leftarrow \emptyset$. For for every subset $X \subseteq \mathsf{F} \setminus \mathsf{L}$ of size at most $b+1$, we verify that X by itself is independent, and furthermore, that the set $Y \subseteq \mathsf{L}$ of objects intersecting the objects of X, is of size at most $|X|-1$. If so, we do $\mathsf{L} \leftarrow (\mathsf{L} \setminus Y) \cup X$. Every such exchange increases the size of L by at least one, and as such it can happen at most n times. Naively, there are $\binom{n}{b+1}$ subsets X to consider, and for each such subset X it takes $O(nb)$ time to compute Y. As such the running time is bounded by $O(n^{b+3})$. (The running time can be probably improved by being a bit more careful about the implementation.)

3.2 Analysis

We present two alternative ways to analyze this algorithm. The first approach uses only the fact that the union complexity is low. The second approach is more direct, and uses the property that the regions are admissible.

3.2.1 Analysis using union complexity

The following lemma by Afshani and Chan [AC06], which was originally intended for different purposes, will turn out to be useful here (the proof exploits linearity of planar graphs and the Clarkson technique [CS89]):

Lemma 3.1 *Suppose we have n disjoint simply connected regions in the plane and a collection of disjoint curves, where each curve intersects at most k regions. Call two curves* equivalent *if they intersect precisely the same subset of regions. Then the number of equivalent classes is at most $c_0 n k^2$ for some constant c_0.*

Let $\mathcal{S}$ be an optimal solution, and let L be a b-locally optimal solution. We will upper-bound $|\mathcal{S}|$ in terms of $|\mathsf{L}|$.

Let $\mathcal{S}_{>b}$ denote the set of objects in $\mathcal{S}$ that intersect at least $b+1$ objects of L. Let $\mathcal{S}_{\leq b}$ be the set of remaining objects in $\mathcal{S}$.

If $f_i \in \mathcal{S}$ intersects $f_j \in \mathsf{L}$, then the pair of objects contributes at least two vertices to the boundary of the union of $\mathcal{S} \cup \mathsf{L}$. Indeed, the objects of $\mathcal{S}$ (resp. L) are disjoint from each other since this is an independent set, and no object is contained inside another (by assumption). Thus,

$$2(b+1)|\mathcal{S}_{>b}| \le \varrho\left(|\mathcal{S}_{>b}| + |\mathsf{L}|\right) \quad \Longrightarrow \quad |\mathcal{S}_{>b}| \le \frac{\varrho}{2(b+1)-\varrho}|\mathsf{L}|.$$

On the other hand, by applying Lemma 3.1 with L as the regions and the boundaries of $\mathcal{S}_{\le b}$ as the curves, the objects in $\mathcal{S}_{\le b}$ form at most $c_0 b^2|\mathsf{L}|$ equivalent classes. Each equivalent class contains at most b objects: Otherwise we would be able to remove b objects from L and intersect $b+1$ objects in this equivalence class to get an independent set larger than L. This would contradict the b-local optimality of L. Thus, $|\mathcal{S}_{\le b}| \le c_0 b^3|\mathsf{L}|$.

Combining the two inequalities, we get

$$|\mathcal{S}| \le |\mathcal{S}_{\le b}| + |\mathcal{S}_{>b}| \le \left(c_0 b^3 + \frac{\varrho}{2(b+1)-\varrho}\right)|\mathsf{L}|.$$

For example, we can set $b = \lceil \varrho/2 \rceil$ and the approximation factor is $O(\varrho^3)$.

Theorem 3.2 *Given a set of n unweighted objects in the plane with linear union complexity, for a sufficiently large constant b, any b-locally optimal independent set has size $\Omega(\mathrm{opt})$, where* opt *is the size of the maximum independent set of the objects.*

3.2.2 Better analysis for admissible regions

A set of regions F is ***admissible***, if for any two regions $f, f' \in \mathsf{F}$, we have that $f \setminus f'$ and $f' \setminus f$ are both simply connected (i.e., connected and contains no holes). Note, that we do not care how many times the boundaries of the two regions intersects, and furthermore, by definition, no region is contained inside another.

Lemma 3.3 *Let F be a set of admissible regions, and consider a independent set of regions $I \subseteq \mathsf{F}$, and a region $f \in \mathsf{F} \setminus I$. Then, the* **core** *region $f \setminus I = f \setminus \bigcup_{g \in I} g$ is non-empty and simply connected.*

Proof: It is easy to verify that for the regions of I to split f into two connected components, they must intersect, which contradicts their disjointness. ■

Lemma 3.4 *Let $X, Y \subseteq \mathsf{F}$ be two independent sets of regions. Then the intersection graph G of $X \cup Y$ is planar.*

Proof: The planarity of this graph follows from Lemma 3.3.

Indeed, for a region $f \in X$, the core $f' = f \setminus \bigcup_{g \in Y} g$ is non-empty and simply connected. Place a vertex v_f inside this region, and for every object $g \in Y$ that intersects f, create a curve from v_f to a point $p_{f,g}$ on the boundary of g that lies inside f. Clearly, we can create these curves in such a way that they do not intersect each other. See figure on the right.

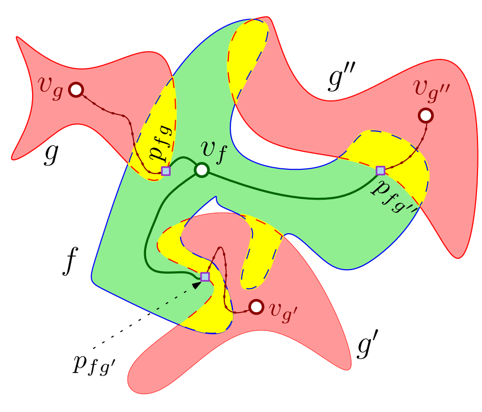

Similarly, for every region $g \in Y$, we place a vertex v_g inside g, and connect it to all the points $p_{f,g}$ placed on its boundary, by curves that are contained in g, and they are interior disjoint. Clearly, together, these vertices and curves form a planar drawing of G. ■

We need the following version of the planar separator theorem. Below, for a set of vertices U in a graph G, let $\Gamma(U)$ denote the set of neighbors of U, and let $\overline{\Gamma}(U) = \Gamma(U) \cup U$.

Lemma 3.5 ([Fre87]) *There are constants c_1, c_2 and c_3, such that for any planar graph $G = (V, E)$ with n vertices, and a parameter r, then one can find a set of $X \subseteq V$ of size at most $c_1 n/\sqrt{r}$, and a partition of $V \setminus X$ into n/r sets $V_1, \ldots, V_{n/r}$, satisfying: (i) $|V_i| \leq c_2 r$, (ii) $\Gamma(V_i) \cap V_j = \emptyset$, for $i \neq j$, and (iii) $|\Gamma(V_i) \cap X| \leq c_3\sqrt{r}$.*

Let $\mathcal{S}$ be the optimal solution and L be a b-locally optimal solution. Consider the bipartite intersection graph G of $\mathcal{S} \cup \mathsf{L}$. By Lemma 3.4, we can apply Lemma 3.5 to G, for $r = b/(c_2+c_3)$. Note that $\left|\overline{\Gamma}(V_i)\right| \leq c_2 r + c_3\sqrt{r} < b$ for each i. Let

$$s_i = |V_i \cap \mathcal{S}|, \quad \ell_i = |V_i \cap \mathsf{L}|, \quad \text{and} \quad b_i = |\Gamma(V_i) \cap X|, \quad \text{for each } i.$$

Observe that $\ell_i + b_i \geq s_i$, for all i. Indeed, otherwise, we can throw away the vertices of $\mathsf{L} \cap \overline{\Gamma}(V_i)$ from L, and replace them by $V_i \cap \mathcal{S}$, resulting in a better solution. This would contradict the local optimality of L. Thus,

$$\begin{aligned} |\mathcal{S}| \;&\leq\; \sum_i s_i \;+\; |X| \;\leq\; \sum_i \ell_i \;+\; \sum_i b_i \;+\; |X| \;\leq\; |\mathsf{L}| + c_2\sqrt{r}\cdot\frac{|\mathcal{S}|+|\mathsf{L}|}{r} + c_1\frac{|\mathcal{S}|+|\mathsf{L}|}{\sqrt{r}} \\ &\leq\; |\mathsf{L}| + (c_1+c_2)\frac{|\mathcal{S}|+|\mathsf{L}|}{\sqrt{r}}. \end{aligned}$$

It follows that $|\mathcal{S}| \leq (1 + O(1/\sqrt{b}))|\mathsf{L}|$. We can set b to the order of $1/\varepsilon^2$, and we get the following.

Theorem 3.6 *Given a set of n unweighted admissible regions in the plane, any b-locally optimal independent set has size $\geq (1 - O(1/\sqrt{b}))\mathrm{opt}$, where opt is the size of the maximum independent set of the objects. In particular, one can compute an independent set of size $\geq (1-\varepsilon)\mathrm{opt}$, in time $n^{O(1/\varepsilon^2)}$.*

4 Approximation using LP relaxation

4.1 The algorithm

We are interested in computing a maximum-weight independent set of the objects in $\mathsf{F} = \{f_1, \ldots, f_n\}$, with weights $w_1, \ldots, w_n$, respectively. To this end, let us solve the following LP relaxation.

$$\begin{aligned} \max \quad & \sum_{i=1}^{n} w_i x_i \\ & \sum_{\mathsf{p} \in f_i} x_i \leq 1 & \qquad \forall \mathsf{p} \in \mathcal{V}(S) \\ & 0 \leq x_i \leq 1, \end{aligned}$$

where $\mathcal{V}(\mathsf{F})$ denotes the set of vertices of the arrangement $\mathcal{A}(\mathsf{F})$.

In the following, x_i will refer to the value assigned to this variable by the solution of the LP. Similarly, $\mathrm{Opt} = \sum_i w_i x_i$ will denote the weight of the relaxed optimal solution, which is at least the weight opt of the optimal integral solution.

We will assume, for the time being, that no two objects of F fully contain each other.

In the unweighted case, our strategy is simple: Randomly put each object f_i into R with probability x_i, for $i = 1, \ldots, n$. Then apply Turán's theorem to output an independent set T of size at least $|\mathsf{R}|/(\Delta + 1)$, where Δ denotes the average degree in the intersection graph formed by R.

In the weighted case, we use a different strategy: Randomly put each object f_i into R with probability x_i/α, for $i = 1, \ldots, n$, for a suitable constant $\alpha > 1$. We extract an independent set out of R greedily as follows. Arrange the objects in decreasing order of weight, so that $w_1 \geq w_2 \geq \cdots \geq w_n$. Initially, we set $\mathsf{T} \leftarrow \emptyset$. Now, we scan the objects of R in order. In the ith step, if $f_i \in \mathsf{R}$ and f_i does not intersect any other object in T, then we insert it into T. Clearly, the objects in T are independent, and we output this set.

4.2 The analysis

4.2.1 Unweighted case

Let

$$\mathcal{E} = \mathcal{E}(\mathsf{F}) = \sum_i x_i.$$

We treat a vertex p of $\mathcal{A}(\mathsf{F})$ as a triple $(\mathsf{p}, i, j) \in \mathcal{V}(\mathsf{F})$, where p is formed by the intersection of the boundaries of f_i and f_j.

The key to our analysis lies in the following inequality, which we prove by an interesting adaptation of the Clarkson technique [CS89].

Lemma 4.1 $\displaystyle\sum_{(\mathsf{p},i,j)\in\mathcal{V}(\mathsf{F})} x_i x_j \;\leq\; 4\varrho\mathcal{E}.$

Proof: Consider a random sample R' of F, where an object f_i is being picked up with probability $x_i/2$. Clearly, we have that $(\mathsf{p}, i, j) \in \mathcal{V}(\mathsf{F})$ appears on the boundary of the union of the objects of R, if and only if f_i and f_j are being picked, and none of the objects that cover p are being chosen into R'. In particular, let $\mathcal{U}(\mathsf{R}')$ denote the vertices on the boundary of the union of the objects of R'. We have that

$$\mathbf{Pr}\Big[\,(\mathsf{p}, i, j) \in \mathcal{U}(\mathsf{R}')\Big] \;=\; \frac{x_i}{2}\cdot\frac{x_j}{2}\prod_{\substack{\mathsf{p}\in f_k,\\ k\neq i, k\neq j}}\Big(1 - \frac{x_k}{2}\Big) \;\geq\; \frac{x_i x_j}{4}\left[1 - \sum_{\substack{\mathsf{p}\in f_k,\\ k\neq i, k\neq j}}\frac{x_k}{2}\right] \;\geq\; \frac{x_i x_j}{8},$$

by the inequality ($\prod_k(1 - a_k) \geq 1 - \sum_k a_k$ for $a_k \in [0, 1]$), since $\sum_{\mathsf{p}\in f_k} x_k \leq 1$ (as the LP solution is valid). On the other hand, the expected number of vertices on the union is at most $\varrho\,\mathbf{E}[\sum_i (x_i/2)] = \varrho\mathcal{E}/2$, which implies the result since

$$\sum_{(\mathsf{p},i,j)\in\mathcal{V}(\mathsf{F})}\frac{x_i x_j}{8} \;\leq\; \sum_{(\mathsf{p},i,j)\in\mathcal{V}(\mathsf{F})}\mathbf{Pr}\Big[(\mathsf{p}, i, j) \in \mathcal{U}(\mathsf{R}')\Big] \;=\; \mathbf{E}\Big[\,|\mathcal{U}(\mathsf{R}')|\,\Big] \;\leq\; \frac{\varrho\mathcal{E}}{2}. \qquad \blacksquare$$

To understand the significance of the above lemma, recall that in the unweighted case, we draw a random sample R where f_i is chosen with probability x_i. The expected size of R is given by $\mathcal{E}$. Let K denotes the number of pairs of intersecting objects in R. The expected value of K is at most $(1/2) \sum_{(\mathsf{p},i,j)\in\mathcal{V}(\mathsf{F})} x_i x_j$. By the above lemma,

$$\mathbf{E}[\mathsf{K}] \leq 2\varrho\mathcal{E};$$

in other words, the intersection graph of R has linear expected number of edges in the size of R. This implies that the average degree $\Delta = 2\mathsf{K}/|\mathsf{R}|$ is a constant, and as such this graph has an independent set that is a constant fraction of the size of R, and as such a constant fraction of the (optimal) LP solution. Formally, we have

$$\mathbf{E}[|\mathsf{R}|]^2 \leq \mathbf{E}\left[\frac{|\mathsf{R}|^2}{2\mathsf{K}+|\mathsf{R}|}\right]\mathbf{E}[2\mathsf{K}+|\mathsf{R}|],$$

by the Cauchy–Schwarz inequality stated for expectations (i.e., $\mathbf{E}[XY]^2 \leq \mathbf{E}[X^2]\,\mathbf{E}[Y^2]$). As such,

$$\begin{aligned}\mathbf{E}[|\mathsf{T}|] \;&\geq\; \mathbf{E}\left[\frac{|\mathsf{R}|}{2\mathsf{K}/|\mathsf{R}|+1}\right] \;=\; \mathbf{E}\left[\frac{|\mathsf{R}|^2}{2\mathsf{K}+|\mathsf{R}|}\right] \;\geq\; \frac{\mathbf{E}[|\mathsf{R}|]^2}{\mathbf{E}[2\mathsf{K}+|\mathsf{R}|]} \;=\; \frac{\mathbf{E}[|\mathsf{R}|]}{2\,\mathbf{E}[\mathsf{K}]\,/\,\mathbf{E}[|\mathsf{R}|]+1}\\ &\geq\; \frac{\mathcal{E}}{4\varrho+1} \;=\; \frac{\mathrm{Opt}}{4\varrho+1},\end{aligned}$$

4.2.2 Analysis – weighted case

In the weighted case, the analysis requires additional ideas. First we need to consider prefixes of F. Let $\mathsf{F}_k = \{f_1, \ldots, f_k\}$ contain the k heaviest objects. Let $\mathcal{E}_k = \sum_{i=1}^{k} x_i$ and

$$\rho_k = \sum_{\substack{i\leq k-1,\\ f_i\cap f_k\neq\emptyset}} x_i.$$

Lemma 4.1 applied to F_k implies that

Lemma 4.2 $\displaystyle\sum_{i=1}^{k} \rho_i x_i \;\leq\; 2\varrho\mathcal{E}_k = 2\varrho\sum_{i=1}^{k} x_i.$

Recall that we draw a random sample R where f_i is chosen with probability x_i/α. Ignoring constant factors, the value ρ_i then represents the expected in-degree of f_i in the intersection graph of R, if we direct heavier objects towards lighter objects. Objects f_i with large ρ_i values will be "problematic" since it is unlikely to end up in the generated independent set, as the probability for that is roughly $(x_i/\alpha)\exp(-\rho_i/\alpha)$. However, the above lemma says that (intuitively) the objects in R should have small ρ_i values on average, over any prefix of objects. As such, we can "charge" the problematic objects to heavier regular objects. The next lemma says that objects f_i with small ρ_i values have a good chance of showing up in the output independent set.

Lemma 4.3 $p_i \;:=\; \mathbf{Pr}[f_i \in \mathsf{T}] \;\geq\; \frac{x_i}{\alpha}\Big(1 - \frac{\rho_i}{\alpha}\Big).$

Proof: The probability that f_i survives in T is at least

$$\frac{x_i}{\alpha} \prod_{k<i, f_k \cap f_i \neq \emptyset} \Big(1 - \frac{x_k}{\alpha}\Big) \;\geq\; \frac{x_i}{\alpha}\left[1 - \sum_{k<i, f_k \cap f_i \neq \emptyset} \frac{x_k}{\alpha}\right]$$

by the product-sum inequality mentioned earlier. ■

Combining the two lemmas, we get

$$\sum_{i=1}^{k} p_i \;\geq\; \frac{1}{\alpha}\left[\sum_{i=1}^{k} x_i - \frac{1}{\alpha}\sum_{i=1}^{k} \rho_i x_i\right] \;\geq\; \frac{1}{\alpha}\Big(1 - \frac{2\varrho}{\alpha}\Big)\sum_{i=1}^{k} x_i \;\geq\; \frac{1}{8\varrho}\sum_{i=1}^{k} x_i$$

by setting $\alpha = 4\varrho$. As the above holds for all k and the w_i's are in decreasing order, we can conclude that

$$\mathbf{E}\left[\sum_{f_i \in \mathsf{T}} w_i\right] \;=\; \sum_{i=1}^{n} w_i p_i \;\geq\; \frac{1}{8\varrho}\sum_{i=1}^{n} w_i x_i \;=\; \frac{\mathrm{Opt}}{8\varrho},$$

by setting $u_i = p_i$ and $v_i = x_i/8\varrho$ in the following easy technical lemma.

Lemma 4.4 *Let* $u_1, \ldots, u_n$*,* $v_1, \ldots, v_n$ *be non-negative numbers, such that for any* k*, we have that* $\sum_{i=1}^{k} u_i \geq \sum_{i=1}^{k} v_i$*. Let* $w_1 \geq w_2 \geq \cdots \geq w_n \geq 0$ *be real numbers. Then, we have that* $\sum_{i=1}^{n} w_i u_i \geq \sum_{i=1}^{n} w_i v_i$*.*

Proof: Indeed, setting $w_{n+1} = 0$, we have that

$$\sum_{i=1}^{n} w_i u_i = \sum_{i=1}^{n}\left[(w_i - w_{i+1})\sum_{k=1}^{i} u_k\right] \geq \sum_{i=1}^{n}\left[(w_i - w_{i+1})\sum_{k=1}^{i} v_k\right] = \sum_{i=1}^{n} w_i v_i,$$

since $w_i - w_{i+1} \geq 0$, for all i. ■

4.3 Remarks

Derandomization. The variance of $\sum_{f_i \in \mathsf{T}} w_i$ could be high, but fortunately the algorithm can be derandomized by the standard method of conditional probabilities/expectations [MR95]. To this end, it would be simpler to work with a modified definition of T: put f_i in T if and only if $f_i \in \mathsf{R}$ and $f_j \notin \mathsf{R}$ for all $j < i$ with $f_i \cap f_j \neq \emptyset$. Lemma 4.3 still holds, and so the analysis still goes through, but the advantage is that we can now calculate the exact value of $\mathbf{E}\Big[\sum_{f_i \in \mathsf{T}} w_i\Big]$ easily (in polynomial time), even when conditioned to the events that some objects are known to be in or not in T. We can thus deterministically examine each object one by one and determine whether to place it in T by calculating various conditional expectations.

Coping with object containment. We have assumed that no object is fully contained in another, but this assumption can be removed by adding the constraint $\sum_{f_i \subset f_j} x_j \leq 1$ for each i to the LP. Then clearly we have $\sum_{f_i \subset f_j,\ i,j \leq k} x_i x_j \leq \mathcal{E}_k$, and so Lemma 4.2 modifies to $\sum_{i=1}^{k} \rho_i x_i \leq (2\varrho + 1)\mathcal{E}_k$. The approximation factor then readjusts to $4(2\varrho + 1)$.

Time to solve the LP. This LP is a packing LP with $O(n^2)$ inequalities, and n variables. As such, it can be $(1+\varepsilon)$-approximated in $O(n^3 + \varepsilon^{-2} n^2 \log n) = O(n^3)$ by a randomized algorithm that succeeds with high probability [KY07]. To our purposes, it is sufficient to set ε to be a sufficient small constant, say $\varepsilon = 10^{-4}$.

We have thus proved:

Theorem 4.5 *Given a set of n weighted objects in the plane with linear union complexity, one can compute in polynomial time an independent set of total weight* $\Omega(\mathrm{opt})$*, where* opt *is the maximum weight over all independent sets of the objects. The running time of the algorithm is $O(n^3)$.*

A combinatorial result. In the unweighted case, we obtain the following result as a byproduct:

Theorem 4.6 *Given a set of n pseudo-disks in the plane, let* opt *be the size of the maximum independent set and let* opt' *be the size of the minimum set of points that pierce all the pseudo-disks. Then* $\mathrm{opt} = \Omega(\mathrm{opt}')$.

Proof: By the preceding analysis, we have $\mathrm{opt} = \Omega(\mathrm{Opt})$, i.e., the integrality gap of our LP is a constant.

For piercing, the LP relaxation is

$$
\begin{aligned}
\min \quad & \sum_{\mathsf{p} \in \mathcal{V}(S)}^{n} y_{\mathsf{p}} & \\
& \sum_{\mathsf{p} \in f_i} y_{\mathsf{p}} \geq 1 & \forall i = 1, \ldots, n \\
& 0 \leq y_{\mathsf{p}} \leq 1. &
\end{aligned}
$$

Let Opt' be the value of this LP. Known analysis [Lon01, ERS05] implies that the integrality gap of this LP is constant if there exist ε-nets of linear size for a corresponding class of hypergraphs formed by objects in F and points in $\mathcal{V}(S)$. Pyrga and Ray [PR08, Theorem 12] obtained such an existence proof for this ("primal") hypergraph for pseudo-disks. Thus, $\mathrm{opt}' = O(\mathrm{Opt}')$.

To conclude, observe that the two LPs are precisely the dual of each other, and so $\mathrm{Opt} = \mathrm{Opt}'$. ∎

References

[AC06] P. Afshani and T. M. Chan. Dynamic connectivity for axis-parallel rectangles. In *Proc. 14th European Sympos. Algorithms*, volume 4168 of *Lecture Notes Comput. Sci.*, pages 16–27, 2006.

[AG60] E. Asplund and B. Grübaum. On a coloring problem. *Math. Scand.*, 8:181–188, 1960.

[AM06] P. K. Agarwal and N. H. Mustafa. Independent set of intersection graphs of convex objects in 2D. *Comput. Geom. Theory Appl.*, 34(2):83–95, 2006.

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

[Aro98] S. Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. *J. Assoc. Comput. Mach.*, 45(5):753–782, Sep 1998.

[AS00] N. Alon and J. H. Spencer. *The Probabilistic Method*. Wiley Inter-Science, 2nd edition, 2000.

[AvKS98] P. K. Agarwal, M. van Kreveld, and S. Suri. Label placement by maximum independent set in rectangles. *Comput. Geom. Theory Appl.*, 11:209–218, 1998.

[Bak94] B. S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. Assoc. Comput. Mach.*, 41:153–180, 1994.

[BDMR01] P. Berman, B. DasGupta, S. Muthukrishnan, and S. Ramaswami. Efficient approximation algorithms for tiling and packing problems with rectangles. *J. Algorithms*, 41:443–470, 2001.

[BF99] P. Berman and T. Fujito. On approximation properties of the independent set problem for low degree graphs. *Theo. Comp. Sci.*, 32(2):115–132, 1999.

[BG95] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete Comput. Geom.*, 14:263–279, 1995.

[CC09] P. Chalermsook and J. Chuzhoy. Maximum independent set of rectangles. In *Proc. 20th ACM-SIAM Sympos. Discrete Algorithms*, 2009. to appear.

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

[Cha04] T. M. Chan. A note on maximum independent sets in rectangle intersection graphs. *Inform. Process. Lett.*, 89:19–23, 2004.

[CS89] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

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

[EJS05] T. Erlebach, K. Jansen, and E. Seidel. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.*, 34(6):1302–1323, 2005.

[EKNS00] A. Efrat, M. J. Katz, F. Nielsen, and M. Sharir. Dynamic data structures for fat objects and their applications. *Comput. Geom. Theory Appl.*, 15:215–227, 2000.

[ERS05] G. Even, D. Rawitz, and S. Shahar. Hitting sets when the VC-dimension is small. *Inform. Process. Lett.*, 95(2):358–362, 2005.

[Fre87] G. N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.*, 16(6):1004–1022, 1987.

[Hal98] M. M. Halldórsson. Approximations of independent sets in graphs. In *The 2nd Intl. Work. Approx. Algs. Combin. Opt. Problems*, pages 1–13, 1998.

[Has96] J. Hastad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, pages 627–636, 1996.

[KLPS86] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.*, 1:59–71, 1986.

[KY07] C. Koufogiannakis and N. E. Young. Beating simplex for fractional packing and covering linear programs. In *Proc. 48th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 494–506, 2007.

[Lon01] P. M. Long. Using the pseudo-dimension to analyze approximation algorithms for integer programming. In *Proc. 7th Workshop Algorithms Data Struct.*, volume 2125 of *Lecture Notes Comput. Sci.*, pages 26–37, 2001.

[LT79] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.

[MBH+95] M. V. Marathe, H. Breu, H. B. Hunt III, S. S. Ravi, and D. J. Rosenkrantz. Simple heuristics for unit disk graphs. *Networks*, 25:59–68, 1995.

[MR95] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge University Press, New York, NY, 1995.

[NHK05] T. Nieberg, J. Hurink, and W. Kern. A robust ptas for maximum weight independent set in unit disk graphs. In *Proc. 30th Int. Workshop Graph-Theoretic Concepts in Comput. Sci.*, volume 3353 of *Lecture Notes Comput. Sci.*, pages 214–221, 2005.

[PR08] E. Pyrga and S. Ray. New existence proofs ε-nets. In *Proc. 24th ACM Sympos. Comput. Geom.*, pages 199–207, 2008.

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

[WZ90] S. Whitesides and R. Zhao. k-admissible collections of Jordan curves and offsets of circular arc figures. Technical Report SOCS 90.08, McGill Univ., Montreal, PQ, 1990.

A Weighted Rectangles

A.1 The algorithm

For the case of weighted axis-aligned rectangles, we can solve the same LP, where the set $\mathcal{V}(\mathsf{F})$ contains both intersection points and corners of the given rectangles.

As before, randomly put each object f_i into R with probability x_i/b, for $i = 1, \ldots, n$, for a suitable constant $b > 1$.

Define two subgraphs $\mathcal{G}_1$ and $\mathcal{G}_2$ of the intersection graph of R: if the boundaries of f_i and f_j intersect zero or two times, put $f_i f_j$ in $\mathcal{G}_1$; if the boundaries intersect four times instead, put $f_i f_j$ in $\mathcal{G}_2$.

We first extract an independent set of $\mathcal{G}_1$ greedily, as before: Arrange the objects so that $w_1 \geq w_2 \geq \cdots \geq w_n$. Put $f_i \in \mathsf{T}$ if and only if $f_i \in \mathsf{R}$ and $f_j \notin \mathsf{R}$ for all $j < i$ with $f_i \cap f_j \neq \emptyset$.

It is well known (e.g., see [AG60]) that $\mathcal{G}_2$ forms a perfect graph (specifically, a comparability graph), so find a Δ-coloring of the rectangles of R in $\mathcal{G}_2$, where Δ denotes the maximum clique size, i.e., the maximum depth in $\mathcal{A}(\mathsf{R})$. Let T' be the color subclass of T of the largest total weight. Clearly, the objects in T' are independent, and we output this set.

A.2 The analysis

As in Section 4.2.2, let $\mathcal{E}_k = \sum_{i=1}^{k} x_i$. Let

$$\rho_k = \sum_{\substack{i \leq k-1, \\ f_i f_k \in \mathcal{G}_1}} x_i.$$

Observe that if $f_i f_j \in \mathcal{G}_1$, then f_i contains a corner of f_j or vice versa. Letting V_j denote the corners of f_j, we have

$$\sum_{i=1}^{k} \rho_i x_i \;\leq\; \sum_{j=1}^{k} \sum_{\mathsf{p} \in V_j} \sum_{\mathsf{p} \in f_i} x_i x_j \;\leq\; \sum_{j=1}^{k} \sum_{\mathsf{p} \in V_j} x_j \;=\; 4\mathcal{E}_k.$$

So, Lemma 4.2 still holds (with ϱ replaced by 2). Lemma 4.3 also holds, and the same arguments imply

$$\mathbf{E}\left[\sum_{f_i \in \mathsf{T}} w_i\right] \;\geq\; \Omega(\mathrm{Opt}).$$

To analyze T', we need a new lemma which bounds the maximum depth of R:

Lemma A.1 $\Delta = O(\log n / \log\log n)$ *with probability at least* $1 - 1/n$.

Proof: Fix a parameter $t > 1$. Fix a point $\mathsf{p} \in \mathcal{V}(\mathsf{F})$. The depth of p in $\mathcal{A}(\mathsf{R})$, denoted by $\mathrm{depth}(\mathsf{p}, \mathsf{R})$, is a sum of independent 0-1 random variables with overall mean $\mu = \sum_{\mathsf{p} \in f_i} x_i \leq 1$. By the Chernoff bound [MR95, page 68],

$$\mathbf{Pr}[\mathrm{depth}(\mathsf{p}, \mathsf{R}) > (1+\delta)\mu] \; < \; \left[\frac{e^\delta}{(1+\delta)^{1+\delta}}\right]^\mu$$

for any $\delta > 0$ (possibly large). By setting δ so that $t = (1+\delta)\mu$, this probability becomes less than $(e/t)^t$. Since $|\mathcal{V}(\mathsf{F})| = O(n^2)$, the probability that $\Delta > t$ is at most $O((e/t)^t n^2)$, which is at most $1/n$ by choosing $t = \Theta(\log n / \log\log n)$. ■

By construction of T', we know that

$$\sum_{f_i \in \mathsf{T}'} w_i \; \geq \; \frac{1}{\Delta} \sum_{f_i \in \mathsf{T}} w_i \; \geq \; \frac{1}{t} \sum_{f_i \in \mathsf{T}} w_i - \mathrm{Opt} \cdot 1_{\Delta > t}$$

where 1_A denotes the indicator variable for event A. With $t = \Theta(\log n / \log\log n)$, we conclude that

$$\mathbf{E}\left[\sum_{f_i \in \mathsf{T}'} w_i\right] \; \geq \; \Omega(\log\log n / \log n) \sum_{f_i \in \mathsf{T}'} w_i - \mathrm{Opt}/n \; \geq \; \Omega(\log\log n / \log n) \cdot \mathrm{Opt}.$$

A.3 Remarks

Derandomization. This algorithm can also be derandomized by the method of conditional expectations. The trick is to consider the following random variable

$$Z \; := \; \frac{1}{t} \sum_{f_i \in \mathsf{T}} w_i - \mathrm{Opt} \cdot \sum_{\mathsf{p} \in \mathcal{V}(\mathsf{F})} (1 + \delta_{\mathsf{p}})^{\mathrm{depth}(\mathsf{p}, \mathsf{R}) - t},$$

where δ_{p} is the δ from the proof of Lemma A.1 and t is the same as before. This variable Z lower-bounds $\sum_{f_i \in \mathsf{T}'} w_i$ (the bound is trivially true if $\Delta > t$, since Z would be negative). Our analysis still implies that $\mathbf{E}[Z] \geq \Omega(\log\log n / \log n) \cdot \mathrm{Opt}$ (the standard proof of the Chernoff bound actually shows that $\mathbf{E}\big[(1+\delta_{\mathsf{p}})^{\mathrm{depth}(\mathsf{p}, \mathsf{R}) - t}\big] < (e/t)^t$). The advantage of working with Z is that we can calculate $\mathbf{E}[Z]$ exactly in polynomial time, even when conditioned to the events that some objects are known to be in or not in T (since $\mathrm{depth}(\mathsf{p}, \mathsf{R})$ is a sum of independent 0-1 random variables).

We have thus proved:

Theorem A.2 *Given a set of n weighted axis-aligned boxes in the plane, one can compute in polynomial time an independent set of total weight* $\Omega(\log\log n / \log n) \cdot \mathrm{opt}$, *where* opt *is the maximum weight over all independent sets of the objects.*

Higher dimensions. By a standard divide-and-conquer method [AvKS98], we get an approximation factor of $O(\log^{d-1} n / \log\log n)$ for weighted axis-aligned boxes in any constant dimension d.