Clique-Based Separators for Geometric Intersection Graphs

Mark de Berg
Department of Computer Science, TU Eindhoven, The Netherlands

Sándor Kisfaludi-Bak[1]
Institute for Theoretical Studies, ETH Zürich, Switzerland

Morteza Monemizadeh
Department of Computer Science, TU Eindhoven, The Netherlands

Leonidas Theocharous
Department of Computer Science, TU Eindhoven, The Netherlands

Abstract

Let F be a set of n objects in the plane and let $\mathcal{G}^{\times}(F)$ be its intersection graph. A balanced clique-based separator of $\mathcal{G}^{\times}(F)$ is a set $\mathcal{S}$ consisting of cliques whose removal partitions $\mathcal{G}^{\times}(F)$ into components of size at most δn, for some fixed constant $\delta < 1$. The weight of a clique-based separator is defined as $\sum_{C\in\mathcal{S}} \log(|C|+1)$. Recently De Berg et al. (SICOMP 2020) proved that if S consists of convex fat objects, then $\mathcal{G}^{\times}(F)$ admits a balanced clique-based separator of weight $O(\sqrt{n})$. We extend this result in several directions, obtaining the following results.

- Map graphs admit a balanced clique-based separator of weight $O(\sqrt{n})$, which is tight in the worst case.
- Intersection graphs of pseudo-disks admit a balanced clique-based separator of weight $O(n^{2/3}\log n)$. If the pseudo-disks are polygonal and of total complexity $O(n)$ then the weight of the separator improves to $O(\sqrt{n}\log n)$.
- Intersection graphs of geodesic disks inside a simple polygon admit a balanced clique-based separator of weight $O(n^{2/3}\log n)$.
- Visibility-restricted unit-disk graphs in a polygonal domain with r reflex vertices admit a balanced clique-based separator of weight $O(\sqrt{n} + r\log(n/r))$, which is tight in the worst case.

These results immediately imply sub-exponential algorithms for Maximum Independent Set (and, hence, Vertex Cover), for Feedback Vertex Set, and for q-Coloring for constant q in these graph classes.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Computational geometry, intersection graphs, separator theorems

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2021.22

Related Version *Full Version*: https://arxiv.org/abs/2109.09874

Funding The work in this paper is supported by the Dutch Research Council (NWO) through Gravitation-grant NETWORKS-024.002.003.

[1] The research was conducted while the author was at the Max Planck Institute for Informatics, Saarbrücken, Germany.

32nd International Symposium on Algorithms and Computation (ISAAC 2021).
Editors: Hee-Kap Ahn and Kunihiko Sadakane; Article No. 22; pp. 22:1–22:15
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The famous Planar Separator Theorem states that any planar graph $\mathcal{G} = (V, E)$ with n nodes[2] admits a subset $\mathcal{S} \subset V$ of size $O(\sqrt{n})$ nodes whose removal decomposes $\mathcal{G}$ into connected components of size at most $2n/3$. The subset $\mathcal{S}$ is called a balanced[3] *separator* of $\mathcal{G}$. The theorem was first proved in 1979 by Lipton and Tarjan [19], and it has been instrumental in the design of algorithms for planar graphs: it has been used to design efficient divide-and-conquer algorithms, to design sub-exponential algorithms for various NP-hard graph problems, and to design approximation algorithms for such problems.

The Planar Separator Theorem has been extended to various other graph classes. Our interest lies in *geometric intersection graphs*, where the nodes correspond to geometric objects and there is an arc between two nodes iff the corresponding objects intersect. If the objects are disks, the resulting graph is called a *disk graph*. Disk graphs, and in particular unit-disk graphs, are a popular model for wireless communication networks and have been studied extensively. Miller et al. [22] and Smith and Wormald [27] showed that if F is a set of balls in $\mathbb{R}^d$ of ply at most k – the *ply* of F is the maximum number of objects in F with a common intersection – then the intersection graph of F has a separator of size $O(k^{1/d} n^{1-1/d})$. This was generalized by Chan [4] and Har-Peled and Quanrud [14] to intersection graphs of so-called low-density sets. Separators for string graphs – a string graph is an intersection graph of sets of curves in the plane – have also been considered [13, 18, 21], with Lee [18] showing that a separator of size $O(\sqrt{m})$ exists, where m is the number of arcs of the graph.

Even for simple objects such as disks or squares, one must restrict the ply to obtain a separator of small size. Otherwise the objects can form a single clique, which obviously does not have a separator of sublinear size. To design subexponential algorithms for problems such as Maximum Independent Set, however, one can also work with a separator consisting of a small number of cliques instead of a small number of nodes. Such *clique-based separators* were introduced recently by De Berg et al. [8]. Formally, a clique-based separator of a graph $\mathcal{G}$ is a collection $\mathcal{S}$ of node-disjoint cliques whose union is a balanced separator of $\mathcal{G}$. The *weight* of $\mathcal{S}$ is defined as $\text{weight}(\mathcal{S}) := \sum_{C \in \mathcal{S}} \log(|C| + 1)$. De Berg et al. [8] proved that the intersection graph of any set F of n convex fat objects in the plane admits a clique-based separator of weight $O(\sqrt{n})$, and they used this to obtain algorithms with running time $2^{O(\sqrt{n})}$ for many classic NP-hard problems on such graphs. This running time is optimal, assuming the Exponential-Time Hypothesis (ETH). The result generalizes to convex fat objects in $\mathbb{R}^d$, where the bound on the weight of the clique-based separator becomes $O(n^{1-1/d})$.

The goal of our paper is to investigate whether similar results are possible for non-fat objects in the plane. Note that not all intersection graphs admit clique-based separators of small weight. String graphs, for instance, can have arbitrarily large complete bipartite graphs as induced subgraphs, in which case any balanced clique-based separator has weight $\Omega(n)$.

The first type of intersection graphs we consider are map graphs, which are a natural generalization of planar graphs. The other types are generalizations of disk graphs. One way to generalize disk graphs is to consider fat objects instead of disks, as done by De Berg et al. [8]. We will study three other generalizations, involving non-fat objects: pseudo-disks, geodesic disks, and visibility-restricted unit disks. Next we define the graph classes we consider more precisely; see Fig. 1 for an example of each graph class.

---

[2] We use the terms *node* and *arc* when talking about graphs, and *vertex* and *edge* for geometric objects.

[3] For a separator to be balanced it suffices that the components have size at most δn for some constant $\delta < 1$. When we speak of separators, we always mean balanced separators, unless stated otherwise.

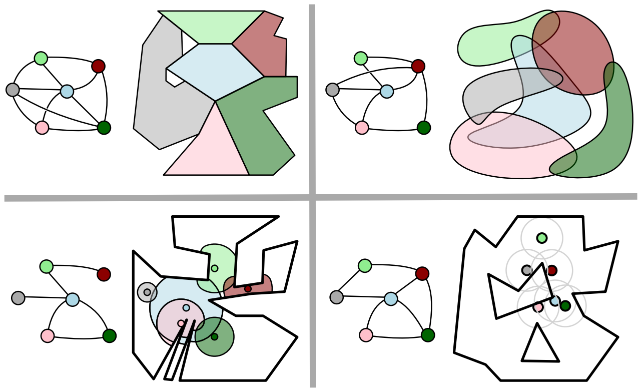

Figure 1 A map graph, a pseudo-disk graph, a geodesic-disk graph, and a visibility restricted unit-disk graph. For the latter class, the grey disks in the picture have radius $\frac{1}{2}$.

In the following, we use $\mathcal{G}^{\times}(F)$ to denote the intersection graph induced by a set F of objects. For convenience, we do not distinguish between the objects and the corresponding nodes, so we use F to denote the set of objects as well as the set of nodes in $\mathcal{G}^{\times}(F)$. We assume that the objects in F are connected, bounded, and closed.

Map graphs. Let $\mathcal{M}$ be a planar subdivision and F be its set of faces. The graph with node set F that has an arc between every pair of neighboring faces is called the *dual graph* of $\mathcal{M}$, and it is planar. Here two faces are neighbors if their boundaries have an edge of the subdivision in common. A *map graph* [6] is defined similarly, except now two faces are neighbors even if their boundaries meet in a single point. Alternatively, we can define a map graph as the intersection graph of a set F of interior-disjoint regions in the plane. Since arbitrarily many faces can share a vertex on their boundary, map graphs can contain arbitrarily large cliques. If at most k faces meet at each subdivision vertex, the graph is called a *k-map graph.* Chen [5] proved that any k-map graph has a (normal, not clique-based) separator of size $O(\sqrt{kn})$, which is also implied by Lee's recent result on string graphs [18].

Pseudo-disk graphs. A set F of objects is a set of pseudo-disks if for any $f, f' \in F$ the boundaries ∂f and $\partial f'$ intersect at most twice. Pseudo-disks were introduced in the context of motion planning by Kedem et al. [15], who proved that the union complexity of n pseudo-disks is $O(n)$. Since then they have been studied extensively. We consider two types of pseudo-disks: polygonal pseudo-disks with $O(n)$ vertices in total, and arbitrary pseudo-disks.

Geodesic-disk graphs and visibility-restricted unit-disk graphs. As mentioned, unit-disk graphs are popular models for wireless communication networks. We consider two natural generalizations of unit-disk graphs, which can be thought of as communication networks in a polygonal environment that may obstruct communication.

- *Geodesic-disk graphs* in a simple polygon P are intersection graphs of geodesic disks inside P. (The *geodesic disk* with center $q \in P$ and radius r is the set of all points in P at geodesic distance at most r from q, where the geodesic distance between two points is the length of the shortest path between them inside P.)

- In *visibility-restricted unit-disk graphs* the nodes correspond to a set Q of n points inside a polygon P, which may have holes, and two points $p, q \in Q$ are connected by an arc iff $|pq| \leqslant 1$ and p and q see each other (meaning that $pq \subset P$).[4] A more general, directed version of such graphs was studied by Ben-Moshe et al. [2] under the name range-restricted visibility graph. They presented an output-sensitive algorithm to compute the graph.

Our results: clique-based separator theorems

So far, clique-based separators were studied for fat objects: De Berg et al. [8] consider convex or similarly-sized fat objects, Kisfaludi-Bak et al. [17] study how the fatness of axis-aligned fat boxes impacts the separator weight, and Kisfaludi-Bak [16] studies balls in hyperbolic space. The $O(\sqrt{n})$ bound on the separator weight is tight even for unit-disk graphs. Indeed, a $\sqrt{n} \times \sqrt{n}$ grid graph can be realized as a unit-disk graph, and any separator of such a grid graph must contain $\Omega(\sqrt{n})$ nodes. Since the maximum clique size in a grid graph is two, any separator must contain $\Omega(\sqrt{n})$ cliques. All graph classes we consider can realize a $\sqrt{n} \times \sqrt{n}$ grid graph, so $\Omega(\sqrt{n})$ is a lower bound on the weight of the clique-based separators we consider. We obtain the following results.

In Section 2 we show that any map graph has a clique-based separator of weight $O(\sqrt{n})$. This gives the first ETH-tight algorithms for Maximum Independent Set (and, hence, Vertex Cover), Feedback Vertex Set, and Coloring in map graphs; see below.

In Section 3 we show that any intersection graph of pseudo-disks has a clique-based separator of weight $O(n^{2/3} \log n)$. If the pseudo-disks are polygonal and of total complexity $O(n)$ then the weight of the separator improves to $O(\sqrt{n} \log n)$.

In Section 4 we consider intersection graphs of geodesic disks inside a simple polygon. At first sight, geodesic disks seem not much harder to deal with than fat objects: they can have skinny parts only in narrow corridors and then packing arguments may still be feasible. Unfortunately another obstacle prevents us from applying a packing argument: geodesic distances in a simply connected polygon induce a metric space whose doubling dimension depends on the number of reflex vertices of the polygon. Nevertheless, by showing that geodesic disks inside a simple polygon behave as pseudo-disks, we are able to obtain a clique-based separator of weight $O(n^{2/3} \log n)$, independent of the number of reflex vertices.

In Section 5 we study visibility-restricted unit-disk graphs. We give an $\Omega(\min(n, r \log(n/r)) + \sqrt{n})$ lower bound for the separator weight, showing that a clique-based separator whose weight depends only on n, the number of points defining the visibility graph, is not possible. We then show how to construct a clique-based separator of weight $O(\min(n, r \log(n/r)) + \sqrt{n})$.

All separators can be computed in polynomial time. For map graphs and for the pseudo-disk intersection graphs, we assume the objects have total complexity $O(n)$. If the objects have curved edges, we assume that basic operations (such as computing the intersection points of two such curves) take $O(1)$ time.

Applications

We apply our separator theorems to obtain subexponential algorithms in the graph classes discussed above, for Maximum Independent Set, Feedback Vertex Set, and q-Coloring for constant q. The crucial property of these problems that makes our separator

[4] Visibility-restricted unit-disk graphs are, strictly speaking, not intersection graphs. In particular, if R_q is defined as the region of points within P that are visible from q and lie within distance $1/2$, then the visibility-restricted unit-disk graph is *not* the same as the intersection graph of the objects R_q.

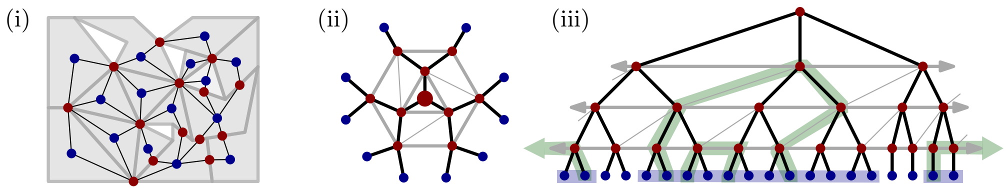

Figure 2 (i) A witness graph for the map graph induced by the grey regions. Points in P are blue, points in Q are red. (ii) The gadget used to replace a witness point. The edges of $\mathcal{T}_q$ are black, the cycles connecting nodes at the same level are grey and thick, the edges to triangulate the 4-cycle are grey and thin. (iii) The green paths show an example of how the separator can intersect a gadget. (Note that the tree "wraps around", as in part (ii) of the figure; see also one of the green paths.) The objects added to the clique C_q correspond to the leaves indicated by the blue rectangles.

applicable, is that the possible ways in which a solution can "interact" with a clique of size k is polynomial in k. We use known techniques (mostly from De Berg et al. [8]) to solve the three problems on any graph class that has small clique-based separators.

All our graph classes are subsumed by string graphs. Bonnet and Rzazewski [3] showed that string graphs have $2^{O(n^{2/3}\log n)}$ algorithms for Maximum Independent Set and 3-Coloring, and a $2^{n^{2/3}\log^{O(1)} n}$ algorithm for Feedback Vertex Set, and that string graphs do not have subexponential algorithms for q-Coloring with $q \geqslant 4$ under ETH. One can also obtain subexponential algorithms in some of our classes from results of Fomin et al. [12, 5] or Marx and Philipczuk [20]. The running times we obtain match or slightly improve the results that can be obtained from these existing results. It should be kept in mind, however, that the existing results are for more general graph classes. An exception are our results on map graphs, which were explicitly studied before and where we improve the running time for Maximum Independent Set and Feedback Vertex Set from $2^{O(\sqrt{n}\log n)}$ to $2^{O(\sqrt{n})}$. (But, admittedly, the existing results apply in the parameterized setting while ours don't.) In any case, the main advantage of our approach is that it allows us to solve Maximum Independent Set, Feedback Vertex Set and q-Coloring on each of the mentioned graph classes in a uniform manner. Refer to the arXiv version [10] for more details on the state of the art, and the specific results we obtain.

2 Map graphs

Recall that a map graph is the intersection graph of a set F of interior-disjoint objects in the plane. We construct a clique-based separator for $\mathcal{G}^\times(F)$ in four steps. First, we construct a bipartite plane *witness graph* $\mathcal{H}_1$ with node set $P \cup Q$, where the nodes in P correspond to the objects in F and the nodes in Q (with their incident arcs) model the adjacencies in $\mathcal{G}^\times(F)$. Next, we replace each node $q \in Q$ by a certain gadget whose "leaves" are the neighbors of q, and we triangulate the resulting graph. We then apply the Planar Separator Theorem to obtain a separator for the resulting graph $\mathcal{H}_2$. Finally, we turn the separator for $\mathcal{H}_2$ into a clique-based separator for $\mathcal{G}^\times(F)$. Next we explain these steps in detail.

Step 1: Creating a witness graph

To construct a witness graph for $\mathcal{G}^\times(F)$ we use the method of Chen et al. [6]: take a point p_f in the interior of each object $f \in F$, and take a *witness point* $q \in \partial f \cap \partial f'$ for each pair of touching objects $f, f' \in F$ and add arcs from q to the points p_f and $p_{f'}$. Let $P = \{p_f : f \in F\}$ and let Q be the set of all witness points added. We denote the resulting bipartite graph with node set $P \cup Q$ by $\mathcal{H}_1$; see Figure 2(i) for an example. Observe that points where many objects meet can serve as witness points for many neighboring pairs in $\mathcal{G}^{\times}(F)$. Chen et al. [6, Lemma 2.3] proved that any map graph admits a witness set Q of size $O(n)$. If the objects in F are polygons with $O(n)$ vertices in total then Q can be found in $O(n)$ time (since the vertices can serve as the set Q.)

Step 2: Replacing witness points by gadgets and triangulating

We would like to construct a separator for $\mathcal{H}_1$ using the Planar Separator Theorem, and convert it to a clique-based separator for $\mathcal{G}^{\times}(F)$. For every witness point $q \in Q$ in the separator for $\mathcal{H}_1$, the conversion would add a clique C_q to the clique-based separator, namely, the clique corresponding to all objects $f \in F$ such that p_f is adjacent to q. However, the node q adds 1 to the separator size, but the clique C_q adds $\log(|C_q|+1)$ to the weight of the clique-based separator. To deal with this we modify $\mathcal{H}_1$, as follows.

Consider a node $q \in Q$. Let $N(q) \subseteq P$ denote the set of neighbors of q. For all nodes $q \in Q$ with $|N(q)| \geqslant 3$, we replace the star induced by $\{q\} \cup N(q)$ by a gadget G_q, which is illustrated in Figure 2(ii) and defined as follows.

First, we create a tree $\mathcal{T}_q$ with root q and whose leaves are the nodes in $N(q)$, as follows. Define the *level* $\ell(v)$ of a node v in $\mathcal{T}_q$ to be the distance of v to the root; thus the root has level 0, its children have level 1, and so on. All leaves in $\mathcal{T}_q$ are at the same level, denoted $\ell_{\max}$. The root has degree 3, nodes at level ℓ with $1 \leqslant \ell < \ell_{\max} - 1$ have degree 2, and nodes at level $\ell_{\max} - 1$ have degree 2 or 1. For each $\ell < \ell_{\max}$ we connect the nodes at level ℓ into a cycle. After doing so, all faces in the gadget (except the outer face) are triangles or 4-cycles. We finish the construction by adding a diagonal in each 4-cycle. Define the *height* of a node v as $\text{height}(v) := \ell_{\max} - \ell(v)$. The following observation follows from the construction.

▶ **Observation 1.** *Let v be a node at height $h > 0$ in the gadget G_q.*
(i) *The subtree of $\mathcal{T}_q$ rooted at v, denoted $\mathcal{T}_q(v)$, has at most $3 \cdot 2^{h-1}$ leaves.*
(ii) *The distance from v to any leaf in $\mathcal{T}_q$ is at least h.*

To unify the exposition, it will be convenient to also create a gadget for the case where q has only two neighbors in $\mathcal{H}_1$, say p_f and $p_{f'}$. We then define $\mathcal{T}_q$ to consist of the arcs (q, p_f) and $(q, p_{f'})$. Note that Observation 1 holds for this gadget as well.

By replacing each witness point $q \in Q$ with a gadget G_q as above, we obtain a (still planar) graph. We triangulate this graph to obtain a maximal planar graph $\mathcal{H}_2$.

Step 3: Constructing a separator for $\mathcal{H}_2$

We now want to apply the Planar Separator Theorem to $\mathcal{H}_2$. Our final goal is to obtain a balanced clique-based separator for $\mathcal{G}^{\times}(F)$. Hence, we want the separator for $\mathcal{H}_2$ to be balanced with respect to P. We will also need the separator for $\mathcal{H}_2$ to be connected. Both properties are guaranteed by the following version of the Planar Separator Theorem, which was proved by Djidjev and Venkatesan [11].

Planar Separator Theorem. Let $\mathcal{G} = (V, E)$ be a maximal planar graph with n nodes. Let each node $v \in V$ have a non-negative cost, denoted $\text{cost}(v)$, with $\sum_{v \in V} \text{cost}(v) = 1$. Then V can be partitioned in $O(n)$ time into three sets $A, B, \mathcal{S}$ such that (i) $\mathcal{S}$ is a simple cycle of size $O(\sqrt{n})$, (ii) $\mathcal{G}$ has no arcs between a node in A and a node in B, and (iii) $\sum_{v \in A} \text{cost}(v) \leqslant 2/3$ and $\sum_{v \in B} \text{cost}(v) \leqslant 2/3$.

When applying the Planar Separator Theorem to $\mathcal{H}_2$, we set $\text{cost}(p) := 1/n$ for all nodes $p_f \in P$ and $\text{cost}(v) := 0$ for all other nodes. We denote the resulting separator for $\mathcal{H}_2$ by $\mathcal{S}(\mathcal{H}_2)$ and the node sets inside and outside the separator by $A(\mathcal{H}_2)$ and $B(\mathcal{H}_2)$, respectively.

Step 4: Turning the separator for $\mathcal{H}_2$ into a clique-based separator for $\mathcal{G}^\times(F)$

We convert $\mathcal{S}(\mathcal{H}_2)$ into a clique-based separator $\mathcal{S}$ for $\mathcal{G}^\times(F)$ as follows.

- For each node $p_f \in \mathcal{S}(\mathcal{H}_2) \cap P$ we put the (singleton) clique $\{f\}$ into $\mathcal{S}$.
- For each gadget G_q we proceed as follows. Let V_q be the set of all nodes $v \in \mathcal{T}_q$ that are in $\mathcal{S}(\mathcal{H}_2)$, and define $C_q := \{f \in F : p_f \text{ is a leaf of } \mathcal{T}_q(v) \text{ that has an ancestor in } V_q\}$; see Figure 2(iii). Observe that C_q is a clique in $\mathcal{G}^\times(F)$. We add[5] C_q to $\mathcal{S}$.

The clique-based separator $\mathcal{S}$ induces a partition of $F \setminus \bigcup_{C\in\mathcal{S}} C$ into two parts A and B, with $|A|, |B| \leqslant 2n/3$, in a natural way, namely as $A := \{f \in F : f \notin \bigcup_{C\in\mathcal{S}} C \text{ and } p_f \in A(\mathcal{H}_2)\}$ and $B := \{f \in F : f \notin \bigcup_{C\in\mathcal{S}} C \text{ and } p_f \in B(\mathcal{H}_2)\}$. The proof that $\mathcal{S}$ is a valid separator, that is, there are no edges between objects in A and B, can be found in the full version [10]. It follows from the fact that for any arc (f, f') with witness q, either an ancestor of p_f or $p_{f'}$ is in V_q (and so f or f' is in the separator) or p_f and $p_{f'}$ (and, hence, f and f') are in the same component after removing the separator. It remains to prove that $\mathcal{S}$ has the desired weight.

▶ **Lemma 2.** *The total weight of the separator $\mathcal{S}$ satisfies $\sum_{C\in\mathcal{S}} \log(|C|+1) = O(\sqrt{n})$.*

Proof. Since $\mathcal{S}(\mathcal{H}_2)$ contains $O(\sqrt{n})$ nodes, it suffices to bound the total weight of the cliques added for the gadgets G_q. Consider a gadget G_q. Recall that V_q is the set of all nodes $v \in \mathcal{T}_q$ that are in $\mathcal{S}(\mathcal{H}_2)$. We claim that $\log(|C_q|+1) = O(|V_q|)$, which implies that $\sum_q \log(|C_q|+1) = \sum_q O(|V_q|) = O(\sqrt{n})$, as desired. It remains to prove the claim.

Since $\mathcal{S}(\mathcal{H}_2)$ is a simple cycle, its intersection with G_q consists of one or more paths. Each path π enters and exits G_q at a node in $N(q)$. Let D_π denote the set of all descendants of the nodes in π. We will prove that $\log(|D_\pi|+1) = O(|\pi|)$, where $|\pi|$ denotes the number of nodes of π. This implies the claim since $\log(|C_q|+1) \leqslant \sum_\pi \log(|D_\pi|+1) = \sum_\pi O(|\pi|) = O(|V_q|)$.

To prove that $\log(|D_\pi|+1) = O(|\pi|)$, let $h_{\max}$ be the maximum height of any node in π. Thus $|\pi| \geqslant h_{\max}$ by Observation 1(ii). Consider all subtrees of height $h_{\max}$ in $\mathcal{T}_q$. If π visits t such subtrees, then $|\pi| \geqslant t$. Moreover, $|D_\pi| \leqslant 3t \cdot 2^{h_{\max}-1}$ by Observation 1(i). Hence, $\log(|D_\pi|+1) \leqslant \log\left(3t \cdot 2^{h_{\max}-1}+1\right) < h_{\max} + \log(3t) = O(\max(h_{\max}, t)) = O(|\pi|)$. ◀

By putting everything together we obtain the following theorem.

▶ **Theorem 3.** *Let F be a set of n interior-disjoint regions in the plane. Then the intersection graph $\mathcal{G}^\times(F)$ has a clique-based balanced separator of weight $O(\sqrt{n})$. The separator can be computed in $O(n)$ time, assuming that the total complexity of the objects in F is $O(n)$.*

3 Pseudo-disk graphs

Our clique-based separator construction for a set F of pseudo-disks uses so-called planar supports, defined as follows. Let $\mathcal{H}$ be a hypergraph with node set Q and hyperedge set H. A graph $\mathcal{G}_{\text{sup}}$ is a *planar support* [26] for $\mathcal{H}$ if $\mathcal{G}_{\text{sup}}$ is a planar graph with node set Q such that for any hyperedge $h \in H$ the subgraph of $\mathcal{G}_{\text{sup}}$ induced by the nodes in h is connected. In our application we let the node set Q correspond to a set of points stabbing all pairwise

[5] We tacitly assume that if an object is in multiple cliques in $\mathcal{S}$, we remove all but one of its occurrences.

intersections between the pseudo-disks, that is, for each intersecting pair $f, f' \in F$ there will be a point $q \in Q$ that lies in $f \cap f'$. The goal is to keep the size of Q small, by capturing all intersecting pairs with few points. The hyperedges are defined by the regions in F, that is, for every $f \in F$ there is a hyperedge $h_f := Q \cap f$. Let $\mathcal{H}_Q(F)$ denote the resulting hypergraph.

▶ **Lemma 4.** *Let F be a set of n objects in the plane, let Q be a set of points stabbing all pairwise intersections in F, and let $\mathcal{H}_Q(F)$ denote the hypergraph as defined above. If $\mathcal{H}_Q(F)$ has a planar support $\mathcal{G}_{\mathrm{sup}}$ then $\mathcal{G}^\times(F)$ has a clique-based separator of size $O(\sqrt{|Q|})$ and weight $O(\sqrt{|Q|}\log n)$.*

Proof. Let $\mathcal{S}(\mathcal{G}_{\mathrm{sup}})$ be a separator for $\mathcal{G}_{\mathrm{sup}}$ of size $O(\sqrt{|Q|})$, which exists by the Planar Separator Theorem, and let $A(\mathcal{G}_{\mathrm{sup}})$ and $B(\mathcal{G}_{\mathrm{sup}})$ be the corresponding separated parts. To ensure an appropriately balanced separator we use the cost-balanced version of the Planar Separator Theorem, as stated in the previous section. For each object $f \in F$ we give one point $q_f \in Q \cap f$ a cost of $1/n$ and all other points cost 0. We call q_f the *representative* of f. (We assume for simplicity that each $f \in F$ intersects at least one other object $f' \in F$, so we can always find a representative. Objects $f \in F$ not intersecting any other object are singletons in $\mathcal{G}^\times(F)$ and can be ignored.) For a point $q \in Q$, define C_q to be the clique in $\mathcal{G}^\times(F)$ consisting of all objects $f \in F$ that contain q. Our clique-based separator $\mathcal{S}$ for $\mathcal{G}^\times(F)$ is now defined as $\mathcal{S} := \{C_q : q \in \mathcal{S}(\mathcal{G}_{\mathrm{sup}})\}$, and the two separated parts are defined as: $A := \{f \in F : f \notin \mathcal{S} \text{ and } q_f \in A(\mathcal{G}_{\mathrm{sup}})\}$ and $B := \{f \in F : f \notin \mathcal{S} \text{ and } q_f \in B(\mathcal{G}_{\mathrm{sup}})\}$. Clearly, the size of $\mathcal{S}$ is $O(\sqrt{|Q|})$ and its weight is $O(\sqrt{|Q|}\log n)$. Moreover, $|A|, |B| \leqslant 2n/3$ because $\mathcal{S}(\mathcal{G}^*)$ is balanced with respect to the node costs.

We claim there are no arcs in $\mathcal{G}^\times(F)$ between a node in A and a node in B. Suppose for a contradiction that there are intersecting objects f, f' such that $f \in A$ and $f' \in B$. By definition of Q there is a point $q \in Q$ that lies in $f \cap f'$. By the planar-support property, the hyperedge h_f induces a connected subgraph of $\mathcal{G}_{\mathrm{sup}}$, so there is a path π that connects q to the representative q_f and such that all nodes of π are points in $f \cap Q$. No node on the path π can be in $\mathcal{S}(\mathcal{G}_{\mathrm{sup}})$, otherwise f is in a clique that was added to $\mathcal{S}$. Similarly, there is a path π' connecting $q_{f'}$ to q such that no point on π' is in $\mathcal{S}(\mathcal{G}_{\mathrm{sup}})$. But then there is a path from q_f to $q_{f'}$ in $\mathcal{G}_{\mathrm{sup}}$ after the removal of $\mathcal{S}(\mathcal{G}_{\mathrm{sup}})$. Hence, q_f and $q_{f'}$ are in the same part of the partition, which contradicts that f and f' are in different parts.

We conclude that $\mathcal{S}$ is a clique-based separator with the desired properties. ◀

▶ Remark 5. The witness set Q in the previous section stabs all pairwise intersections of objects in the map graph, and so $P \cup Q$ stabs all pairwise intersections as well. $P \cap Q$ has planar support, so we can get a separator for map graphs using Lemma 4. Its weight would be $O(\sqrt{n}\log n)$, however, while in the previous section we managed to get $O(\sqrt{n})$ weight.

Polygonal pseudo-disks

We now apply Lemma 4 to obtain a clique-based separator for a set F of polygonal pseudo-disks. To this end, let Q be the set of vertices of the pseudo-disks in F. Observe that whenever two pseudo-disks intersect, one must have a vertex inside the other. Indeed, either one pseudo-disk is entirely inside the other, or an edge e of f intersects an edge e' of f'. In the latter case, one of the two edges ends inside the other pseudo-disk, otherwise there are three intersections between the boundaries. Furthermore, pseudo-disks have the *non-piercing* property: $f \setminus f'$ is connected for any two pseudo-disks f, f'. Raman and Ray [26] proved[6] that the hypergraph $\mathcal{H}_Q(F)$ of a set of non-piercing regions has a planar support for any set Q, so in particular for the set Q just defined. We can thus apply Lemma 4 to compute a clique-based separator for $\mathcal{G}^\times(F)$. The time to compute the separator is dominated by the computation of the planar support, which takes $O(n^3)$ time [26].

▶ **Theorem 6.** *Let F be a set of n polygonal pseudo-disks in the plane with $O(n)$ vertices in total. Then the intersection graph $\mathcal{G}^\times(F)$ has a clique-based balanced separator of size $O(\sqrt{n})$ and weight $O(\sqrt{n}\log n)$, which can be found in $O(n^3)$ time.*

Arbitrary pseudo-disks

To construct a clique-based separator using Lemma 4 we need a small point set Q that stabs all pairwise intersections. Unfortunately, for general pseudo-disks a linear-size set Q that stabs all intersections need not exist: there is a collection of n disks such that stabbing all pairwise intersections requires $\Omega(n^{4/3})$ points. (Such a collection can be derived from a construction with n lines and n points with $\Omega(n^{4/3})$ incidences [23].) Hence, we need some more work before we can apply Lemma 4.

Our separator result for arbitrary pseudo-disks works in a more general setting, namely for sets from a family $\mathcal{F}$ with linear union complexity. (We say that $\mathcal{F}$ has union complexity $U(n)$ if, for any $n \geqslant 1$ and any subset $F \subset \mathcal{F}$ of size n, the union complexity of F is $U(n)$.) Recall that the union complexity of a family of pseudo-disks is $O(n)$ [15]. The next theorem states that such sets admit a clique-based separator of sublinear weight. Note that the bound only depends on the number of objects, not on their complexity.

▶ **Theorem 7.** *Let F be a set of n objects from a family $\mathcal{F}$ of union complexity $U(n)$, where $U(n) \geqslant n$. Then $\mathcal{G}^\times(F)$ has a clique-based separator of size $O((U(n))^{2/3})$ and weight $O((U(n))^{2/3}\log n)$. In particular, if F is a set of pseudo-disks then $\mathcal{G}^\times(F)$ has a clique-based separator of size $O(n^{2/3})$ and weight $O(n^{2/3}\log n)$. The separator can be computed in $O(n^3)$ time, assuming the total complexity of the objects is $O(n)$.*

Proof. We construct the separator $\mathcal{S}$ in two steps.

The first step proceeds as follows. For a point p in the plane, let C_p denote the set of objects from the (current) set F containing p. As long as there is a point p such that $|C_p| > n^{1/3}$, we remove C_p from F and put C_p into $\mathcal{S}$; here n refers to the size of the initial set F. Thus the first step adds $O(n^{2/3})$ cliques to $\mathcal{S}$ with total weight $O(n^{2/3}\log n)$. This step can easily be implemented in $O(n^3)$ time.

In the second step we have a set $F^* \subseteq F$ of n^* objects with ply k, where $n^* \leqslant n$ and $k \leqslant n^{1/3}$. Let $\mathcal{A}(F^*)$ denote the arrangement induced by F^*. Since F^* has ply k, the Clarkson-Shor technique [7] implies that the complexity of the arrangement $\mathcal{A}(F^*)$ is $O(k^2 \cdot U(n^*/k))$. We can compute this arrangement in $O(k^2 \cdot U(n^*/k)\log n) = O(n^2\log n)$ time [9]. Take a point q in each face of the arrangement, and let Q be the resulting set of $O(k^2 \cdot U(n^*/k))$ points. The set Q stabs all pairwise intersections and the dual graph $\mathcal{G}^*$ of the arrangement $\mathcal{A}(F^*)$ is a planar support for the hypergraph $\mathcal{H}_Q(F)$. Hence, by Lemma 4 there is a clique-based separator $\mathcal{S}^*$ for $\mathcal{G}^\times(F)$ of size $O\left(k\sqrt{U\left(n^*/k\right)}\right)$ and

[6] Raman and Ray assume the sets F and Q defining the hypergraph are in general position. Therefore we first slightly perturb the pseudo-disks in F to get them into general position (while keeping the same intersection graph), then we take Q to be a point set coinciding with the vertex set of F, and then we slightly move the points in Q such that the hypergraph remains the same.

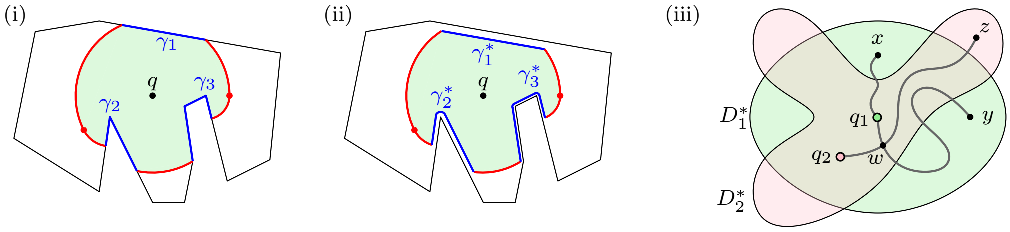

Figure 3 (i) A geodesic disk D with center q and radius r. The set $\Gamma(D)$ has three pieces, γ_1, γ_2 and γ_3, shown in blue. (ii) The result of the perturbation. Note that $|\gamma_1| < |\gamma_2| < |\gamma_3|$ and so $\varepsilon_{\gamma_1} > \varepsilon_{\gamma_2} > \varepsilon_{\gamma_3}$. (iii) Illustration for the proof of Theorem 8.

weight $O\left(k\sqrt{U\left(n^*/k\right)}\log n^*\right)$. Note that $U(n)$ is a superadditive function [1] which implies that $U(n/k) \leqslant U(n)/k$ and therefore $k\sqrt{U\left(n^*/k\right)} \leqslant \sqrt{k\,U(n)} \leqslant (U(n))^{2/3}$. By adding $\mathcal{S}^*$ to the set $\mathcal{S}$ of cliques generated in the first step, we obtain a clique-based separator with the desired properties. ◀

4 Geodesic disks inside a simple polygon

Let P be a simple polygon. We denote the shortest path (or: *geodesic*) in P between two points $p, q \in P$ by $\pi(p, q)$; note that $\pi(p, q)$ is unique since P is simple. The *geodesic distance* between p and q is defined to be $\|\pi(p, q)\|$, where $\|\pi\|$ denotes the Euclidean length of a path π. For a given point $q \in P$ and radius $r > 0$, we call the region $D(q, r) := \{p \in P : \|\pi(p, q)\| \leqslant r\}$ a *geodesic disk*. Let $\mathcal{D} = \{D_1, \ldots, D_n\}$ be a set of geodesic disks in P. To construct a clique-based separator for $\mathcal{G}^\times(\mathcal{D})$ we will show that $\mathcal{D}$ behaves as a set of pseudo-disks so we can apply the result of the previous section.

The structure of a geodesic disk

The boundary $\partial D(q, r)$ of a geodesic disk $D(q, r)$ consists of circular arcs lying in the interior of P (centered at q or at a reflex vertex of P) and parts of the edges of P. We split $\partial D(q, r)$ into *boundary pieces* at the points where the circular arcs meet ∂P. This generates two sets of boundary pieces: a set containing the pieces that consist of circular arcs, and a set $\Gamma(D)$ containing the pieces that consist of parts of edges of P. An example can be seen in Fig. 3.

A region $R \subseteq P$ is *geodesically convex* if for any points $p, q \in R$ we have $\pi(p, q) \subseteq R$. Pollack et al. [25] showed that geodesic disks inside a simple polygon are geodesically convex. An immediate consequence is that the intersection of two geodesic disks is connected.

Geodesic disks behave as pseudo-disks

Geodesic disks in a simple polygon are not proper pseudo-disks. For example, if D_1 and D_2 are the blue and pink pseudo-disk in the third image in Fig. 1, then $D_1 \setminus D_2$ has two components, which is not allowed for pseudo-disks. Nevertheless, we will show that $\mathcal{D}$ behaves as a set of pseudo-disks in the sense that a small perturbation turns them into pseudo-disks, while keeping the intersection graph the same.

As a first step in the perturbation, we increase the radius of each geodesic disk $D_i \in \mathcal{D}$ by some small ε_i. We pick these ε_i such that the intersection graph $\mathcal{G}^\times(\mathcal{D})$ stays the same while all degeneracies disappear. In particular, the boundary pieces of different geodesic disks have different lengths after this perturbation, and no two geodesic disks touch. With a slight abuse of notation, we still denote the resulting set of geodesic disks by $\mathcal{D}$.

The second step in the perturbation moves each $\gamma \in \cup_{i=1}^{n}\Gamma(D_i)$ into the interior of the polygon over some distance ε_γ, which is smaller than any of the perturbation distances chosen in the first step. More formally, for each $\gamma \in \Gamma(D_i)$ we remove all points from D_i that are at distance less than ε_γ from γ; see Fig. 3(ii). To ensure this gives a set of pseudo-disks we choose the perturbation distances ε_γ according to the reverse order of the Euclidean lengths of the pieces. That is, if $\|\gamma\| > \|\gamma'\|$ then we pick $\varepsilon_\gamma < \varepsilon_{\gamma'}$. The crucial property of this scheme is that whenever $\gamma_i \subset \gamma_j$ then γ_i is moved more than γ_j.

We denote the perturbed version of D_i by D_i^* and define $\mathcal{D}^* := \{D_i^* : D_i \in \mathcal{D}\}$. The perturbed versions have the following important property: every connected component of $D_i^* \setminus D_j^*$ contains a point u with $u \in D_i \setminus D_j$, see [10].

▶ **Theorem 8.** *Any set $\mathcal{D}$ of geodesic disks inside a simple polygon P can be slightly perturbed such that the resulting set $\mathcal{D}^*$ is a set of pseudo-disks with $\mathcal{G}^\times(\mathcal{D}) = \mathcal{G}^\times(\mathcal{D}^*)$.*

Proof. Consider the set $\mathcal{D}^*$ resulting from the perturbation described above. Suppose for a contradiction that there exist two objects $D_1^*, D_2^* \in \mathcal{D}^*$ such that ∂D_1^* and ∂D_2^* cross four or more times. Recall that the intersection of two geodesic disks is connected. This property is not invalidated by the perturbation. Hence, if ∂D_1^* and ∂D_2^* cross four or more times then $D_1^* \setminus D_2^*$ (and, similarly, $D_2^* \setminus D_1^*$) has two or more components.

For $i = 1, 2$, let q_i and r_i denote the center and radius of D_i. Without loss of generality assume that $r_1 \leqslant r_2$. Let x and y be points in different components of $D_1^* \setminus D_2^*$; see Fig. 3 (iii). We can pick x and y such that $x, y \in D_1 \setminus D_2$. By concatenating the geodesics $\pi(x, q_1)$ and $\pi(q_1, y)$ we obtain a curve that splits D_2^* into at least two parts – this is independent of where q_1 lies, or whether $\pi(x, q_1)$ and $\pi(q_1, y)$ partially overlap. (Note that these geodesics lie in D_1 but not necessarily in D_1^*. However, they cannot "go around" a component of $D_2^* \setminus D_1^*$, because D_1 cannot fully contain such a component. Hence, $\pi(x, q_1) \cup \pi(q_1, y)$ must indeed go through D_2^*.) Not all components of $D_2^* \setminus D_1^*$ can belong to the same part, otherwise x and y would not be in different components of $D_1^* \setminus D_2^*$. Take a point $z \in D_2^* \setminus D_1^*$ that lies in a different part than q_2, the center of D_2. Again we can pick z such that $z \in D_2 \setminus D_1$. Then the geodesic $\pi(q_2, z)$ must cross $\pi(x, q_1) \cup \pi(q_1, y)$, say at a point $w \in \pi(q_1, y)$. Since $z \notin D_1$ and $y \notin D_2$ we must have $\|\pi(q_1, w) \cup \pi(w, z)\| + \|\pi(q_2, w) \cup \pi(w, y)\| > r_1 + r_2$. But this gives a contradiction because $y \in D_1$ and $z \in D_2$ implies $\|\pi(q_1, w) \cup \pi(w, y)\| + \|\pi(q_2, w) \cup \pi(w, z)\| \leqslant r_1 + r_2$.

It remains to show that $\mathcal{G}^\times(\mathcal{D}) = \mathcal{G}^\times(\mathcal{D}^*)$. As mentioned earlier, the increase of the radii in the first step of the perturbation is chosen sufficiently small so that no new intersections are introduced. The second step shrinks the geodesic disks, so no new intersections are introduced in that step either. Finally, the fact that the perturbations in the second step are smaller than in the first step guarantees that no intersections are removed. ◀

Theorem 8 allows us to apply Theorem 7. When doing so, we actually do not need to perturb the geodesic disks. We only use the perturbation to argue that the number of faces in the arrangement defined by n geodesic disks of ply k is $O(nk)$. Computing the geodesic disks (and then computing the separator) can be done in polynomial time in n and the number of vertices of P. We obtain the following result.

▶ **Corollary 9.** *Let $\mathcal{D}$ be a set of n geodesic disks inside a simple polygon with m vertices. Then $\mathcal{G}^\times(\mathcal{D})$ has a clique-based separator of size $O(n^{2/3})$ and weight $O(n^{2/3} \log n)$, which can be computed in time polynomial in n and m.*

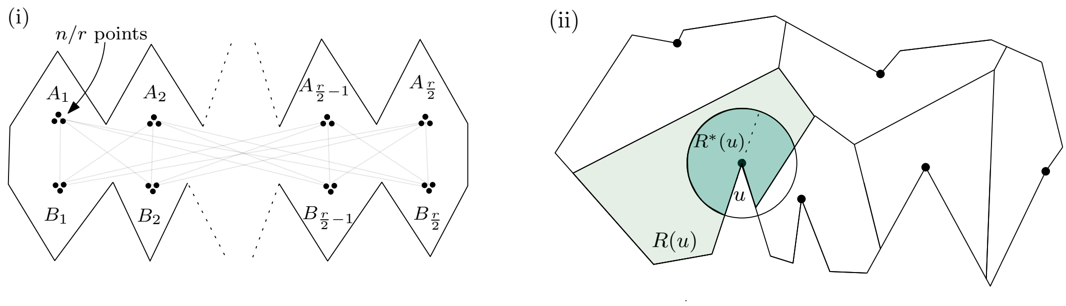

Figure 4 (i) Each cluster A_i sees any of the clusters B_j completely and all distances are at most 1, so a separator that splits $\mathcal{G}^{\times}_{\text{vis},P}\left(\bigcup A_i \cup \bigcup B_j\right)$ into two or more components must fully contain $\bigcup A_i$ or $\bigcup B_j$. Since the clusters A_i (and similarly B_j) do not see each other, such a separator has size at least $r/2$ and weight at least $((r/2)\log(n/r))$. (ii) Splitting $R^*(u)$ into two convex parts.

5 Visibility-restricted unit-disk graphs inside a polygon

Let P be a polygon, possibly with holes, and let Q be a set of n points inside P. We define $\mathcal{G}^{\times}_{\text{vis},P}(Q)$ to be the visibility-restricted unit-disk graph of Q. The nodes in $\mathcal{G}^{\times}_{\text{vis},P}(Q)$ correspond to the points in Q and there is an edge between two points $p, q \in Q$ iff $|pq| \leqslant 1$ and p and q see each other. A vertex of P is *reflex* if its angle within the polygon is more than 180 degrees; note that for a vertex of a hole we look at the angle within P, not within the hole. Below we sketch a proof of the following theorem; a detailed proof is in the arXiv version of the paper [10].

▶ **Theorem 10.** *Let Q be a set of n points inside a polygon (possibly with holes) with r reflex vertices. Then $\mathcal{G}^{\times}_{\text{vis},P}(Q)$ admits a clique-based separator of size $O(\min(n,r)+\sqrt{n})$ and weight $O(\min(n, r\log(n/r)+\sqrt{n}))$. The bounds on the size and weight of the separator are tight in the worst case, even for simple polygons.*

The lower bound

Recall that even for non-visibility restricted unit-disk graphs, $\Omega(\sqrt{n})$ is a lower bound on the worst-case size of the separator. Hence, to prove the lower bound of Theorem 10 it suffices to give an example where the size and weight are $\Omega(\min(n,r))$ and $\Omega(\min(n, r\log(n/r)))$, respectively. This example is given in Fig. 4(i).

The upper bound

We now sketch our separator construction for the case where P is a simple polygon. The extension to polygons with holes can be found in the full version [10].

Step 1: Handling points that see a nearby reflex vertex. Let V_{ref} be the set of reflex vertices of P, and let $Q_1 \subseteq Q$ be the set of points that can see a reflex vertex within distance $\sqrt{2}$. Consider the geodesic Voronoi diagram of V_{ref} within P. Let $R(u)$ be the Voronoi region of vertex $u \in V_{\text{ref}}$ and define $R^*(u) := R(u) \cap D(u, \sqrt{2})$. Note that all points in $R(u)$ can see u and that all points in Q_1 are in $S^*(u)$ for some vertex $u \in V_{\text{ref}}$. By extending one of the edges of P incident to u, we can split $R^*(u)$ into two convex parts; see Fig. 4(ii). Since $R^*(u)$ has diameter $O(1)$, this means that $Q_1 \cap R^*(u)$ can be partitioned into $O(1)$ cliques. We collect all these cliques into a set $\mathcal{S}_1$. Since there are r reflex vertices, $\mathcal{S}_1$ consists of $O(r)$ cliques of total weight is $O(r\log(n/r))$.

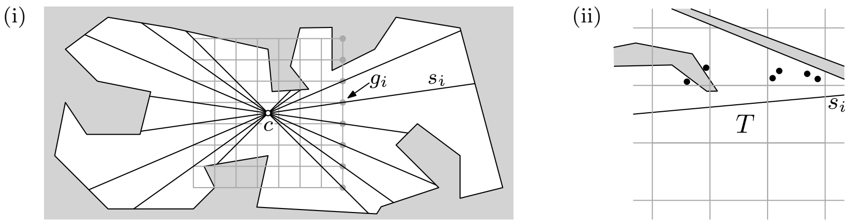

Figure 5 (i) The grid G defining the chords s_i. (ii) The points in the top-left cell see a reflex vertex so they are not in $Q_2(s_i)$. The points in the top-right cell can be split into $O(1)$ cliques.

Step 2: Handling points that do not see a nearby reflex vertex. Our separator $\mathcal{S}$ consists of the cliques in $\mathcal{S}_1$ plus a set $\mathcal{S}_2$ of cliques that are found as follows. Let $Q_2 := Q \setminus Q_1$ be the set of points that do not see a reflex vertex within distance $\sqrt{2}$. Let c be a *centerpoint for Q_2 inside P*, that is, a point such that any (maximal) chord through c splits P into half-polygons containing at most $2|Q_2|/3$ points from Q_2. Such a point always exists; see [10]. Let G be a $\sqrt{n} \times \sqrt{n}$ grid of unit cells centered at c. For each of the $\sqrt{n}$ points g_i in the rightmost column of the grid (even if $g_i \notin P$), we define a chord s_i by taking the line ℓ_i through c and g_i and then taking the component of $\ell_i \cap P$ that contains c; see Fig. 5(i).

For each chord s_i, we define $Q_2(s_i)$ to be the set of points $q \in Q_2$ such that there is a point $z \in s_i$ that sees q with $|qz| \leqslant 1/2$. Note that $\mathcal{G}^{\times}_{\text{vis},P}(Q)$ cannot have an arc between a point $p \in Q_2 \setminus Q_2(s_i)$ above s_i and a point $q \in Q_2 \setminus Q_2(s_i)$ below s_i; otherwise p and/or q see a point on s_i within distance $1/2$, and so at least one of p, q is in $Q_2(s_i)$. Since s_i is a chord through the centerpoint c, this means that s_i induces a balanced separator.

It remains to argue that at least one chord s_i induces a separator of small weight. We will do this by creating a set $\mathcal{S}(s_i)$ of cliques for each chord s_i, and prove that the total weight of these cliques, over all chords s_i, is $O(n)$. Since there are $\sqrt{n}$ chords, one of them has the desired weight.

First, we put all points from $Q_2(s_i)$ that lie outside the grid G into $\mathcal{S}(s_i)$, as singletons. By definition of the chords s_i, a point q outside the grid G lies at distance at most $1/2$ from $O(1)$ chords. Hence, the total number of singleton cliques over all sets $\mathcal{S}(s_i)$ is $O(n)$.

Next, consider a cell T of the grid G. Suppose a point $q \in Q_2(s_i)$ sees a point $z \in T \cap s_i$ with $|qz| \leqslant 1/2$. Then q must lie inside one of the nine grid cells surrounding and including T. Consider such a cell T'. The points in $Q_2(s_i) \cap T'$ that can see a point on $s_i \cap T$ can be partitioned into $O(1)$ cliques; we prove this by showing that if two such points do not see each other, then they must see a reflex vertex within distance $\sqrt{2}$ and, hence, be in Q_1. We thus create $O(1)$ cliques for T' and put them into $\mathcal{S}(s_i)$. This adds at most $O(\log(n_{T'} + 1))$ weight to $\mathcal{S}(s_i)$, where $n_{T'} := |Q_2 \cap T'|$.

Overall, a cell T' adds $O(\log(n_{T'} + 1))$ weight for the chords s_i that cross the nine cells surrounding it. Since there are n cells in total, this immediately gives a total weight of $O(n \log n)$ over all sets $\mathcal{S}(s_i)$. A more careful analysis shows that the total weight of the cliques is actually $O(n)$. Hence, one of the $\sqrt{n}$ chords induces a separator of the desired weight. This finishes the sketch of the construction of the clique-based separator for $\mathcal{G}^{\times}_{\text{vis},P}(Q)$.

6 Concluding Remarks

We showed how clique-based separators with sub-linear weight can be constructed for various classes of intersection graphs which involve non-fat objects. The main advantage of our approach is that we can solve different problems in the graph classes we study in a uniform manner. There are several natural questions that are left open. Some are listed below.

- **Improving the bound for geodesic disks and adding holes.** Our bound on geodesic disks is directly derived by our result on pseudo-disks. However, geodesic disks are much less general than pseudo-disks (and "closer" to regular disks). Hence, one would expect that the optimal weight is closer to $O(\sqrt{n})$. If we allow our polygon to have holes, then our approach for geodesic disks no longer works. Indeed, it is easy to see that even after applying our perturbation scheme, the resulting objects can intersect each other more than two times.
- **Improving the bound for pseudo-disks.** Regarding pseudo-disks, an interesting result [24] states that in every finite family of pseudo-disks in the plane one can find a "small" one, in the sense that it is intersected by only a constant number of disjoint pseudo-disks. This property is also shared by, for instance, convex fat objects. Does this mean that the two graph classes are related in some natural way? If yes, could this connection be exploited to construct separators with better bounds?

References

1. B. Aronov, M.T. Berg, de, E. Ezra, and M. Sharir. Improved bounds for the union of locally fat objects in the plane. *SIAM Journal on Computing*, 43(2):543–572, 2014. doi:10.1137/120891241.
2. Boaz Ben-Moshe, Olaf A. Hall-Holt, Matthew J. Katz, and Joseph S. B. Mitchell. Computing the visibility graph of points within a polygon. In *Proc. 20th ACM Symposium on Computational Geometry*, pages 27–35. ACM, 2004. doi:10.1145/997817.997825.
3. Édouard Bonnet and Pawel Rzazewski. Optimality program in segment and string graphs. *Algorithmica*, 81(7):3047–3073, 2019. doi:10.1007/s00453-019-00568-7.
4. Timothy M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003. doi:10.1016/S0196-6774(02)00294-8.
5. Zhi-Zhong Chen. Approximation algorithms for independent sets in map graphs. *J. Algorithms*, 41(1):20–40, 2001. doi:10.1006/jagm.2001.1178.
6. Zhi-Zhong Chen, Michelangelo Grigni, and Christos H. Papadimitriou. Map graphs. *J. ACM*, 49(2):127–138, 2002. doi:10.1145/506147.506148.
7. Kenneth L. Clarkson and Peter W. Shor. Application of random sampling in computational geometry, II. *Discret. Comput. Geom.*, 4:387–421, 1989. doi:10.1007/BF02187740.
8. Mark de Berg, Hans L. Bodlaender, Sándor Kisfaludi-Bak, Dániel Marx, and Tom C. van der Zanden. A framework for Exponential-Time-Hypothesis-tight algorithms and lower bounds in geometric intersection graphs. *SIAM J. Comput.*, 49:1291–1331, 2020. doi:10.1137/20M1320870.
9. Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. *Computational Geometry: Algorithms and Applications (3rd Edition)*. Springer, 2008. URL: https://www.worldcat.org/oclc/227584184.
10. Mark de Berg, Sándor Kisfaludi-Bak, Morteza Monemizadeh, and Leonidas Theocharous. Clique-based separators for geometric intersection graphs, 2021. arXiv:2109.09874.
11. Hristo Djidjev and Shankar M. Venkatesan. Reduced constants for simple cycle graph separation. *Acta Informatica*, 34(3):231–243, 1997. doi:10.1007/s002360050082.
12. Fedor V. Fomin, Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, and Meirav Zehavi. Decomposition of map graphs with applications. In *Proc. 46th International Colloquium on Automata, Languages, and Programming, (ICALP)*, pages 60:1–60:15, 2019. doi:10.4230/LIPIcs.ICALP.2019.60.
13. Jacob Fox, János Pach, and Csaba D. Tóth. A bipartite strengthening of the crossing lemma. *J. Comb. Theory, Ser. B*, 100(1):23–35, 2010. doi:10.1016/j.jctb.2009.03.005.
14. 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.

15 Klara Kedem, Ron Livne, János Pach, and Micha Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discret. Comput. Geom.*, 1:59–70, 1986. doi:10.1007/BF02187683.

16 Sándor Kisfaludi-Bak. Hyperbolic intersection graphs and (quasi)-polynomial time. In *Proc. 31st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1621–1638, 2020. doi:10.1137/1.9781611975994.100.

17 Sándor Kisfaludi-Bak, Dániel Marx, and Tom C. van der Zanden. How does object fatness impact the complexity of packing in d dimensions? In *30th International Symposium on Algorithms and Computation, ISAAC 2019*, volume 149 of *LIPIcs*, pages 36:1–36:18, 2019. doi:10.4230/LIPIcs.ISAAC.2019.36.

18 James R. Lee. Separators in region intersection graphs. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017*, volume 67 of *LIPIcs*, pages 1:1–1:8, 2017. doi:10.4230/LIPIcs.ITCS.2017.1.

19 Richard J. Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math*, 36(2):177–189, 1977. doi:doi/10.1137/0136016.

20 Dániel Marx and Michal Pilipczuk. Optimal parameterized algorithms for planar facility location problems using voronoi diagrams. In *Proc. 23rd Annual European Symposium on Algorithms (ESA)*, volume 9294 of *Lecture Notes in Computer Science*, pages 865–877. Springer, 2015. doi:10.1007/978-3-662-48350-3_72.

21 Jirí Matoušek. Near-optimal separators in string graphs. *Comb. Probab. Comput.*, 23(1):135–139, 2014. doi:10.1017/S0963548313000400.

22 Gary L. Miller, Shang-Hua Teng, William P. Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. ACM*, 44(1):1–29, 1997. doi:10.1145/256292.256294.

23 János Pach and Micha Sharir. Geometric incidences. In János Pach, editor, *Towards a Theory of Geometric Graphs (Contemporary Mathematics, Vol. 342)*, pages 185–223. Amer. Math. Soc., 2004.

24 Rom Pinchasi. A finite family of pseudodiscs must include a “small” pseudodisc. *SIAM Journal on Discrete Mathematics*, 28:1930–1934, October 2014. doi:10.1137/130949750.

25 Ricky Pollack, Micha Sharir, and Günter Rote. Computing the geodesic center of a simple polygon. *Discret. Comput. Geom.*, 4(6):611–626, 1989. doi:10.1007/BF02187751.

26 Rajiv Raman and Saurabh Ray. Constructing planar support for non-piercing regions. *Discrete & Computational Geometry*, pages 1–25, 2020.

27 Warren D. Smith and Nicholas C. Wormald. Geometric separator theorems & applications. In *39th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 232–243. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743449.