A Clique-Based Separator for Intersection Graphs of Geodesic Disks in $\mathbb{R}^2$

Boris Aronov[1], Mark de Berg[2], and Leonidas Theocharous[3]

1 **Tandon School of Engineering, New York University, Brooklyn NY, USA**
`boris.aronov@nyu.edu`

2 **Department of Mathematics and Computer Science, TU Eindhoven**
`M.T.d.Berg@tue.nl`

3 **Department of Mathematics and Computer Science, TU Eindhoven**
`l.theocharous@tue.nl`

Abstract

A *clique-based separator* of a graph is a (balanced) separator consisting of cliques, and its size is measured as the total number of cliques it contains. Clique-based separators were introduced by De Berg *et al.* [3], who showed that the intersection graph of a set $\mathcal{D}$ of convex fat objects in the plane admits a separator of size $O(\sqrt{n})$. We extend this result by showing that *for any well-behaved shortest-path metric d* defined on a path-connected and closed subset $F \subset \mathbb{R}^2$, a set of geodesic disks with respect to that metric admits a separator consisting of $O(n^{4/5})$ cliques.

1 Introduction

The Planar Separator Theorem states that any planar graph with n nodes has a *balanced separator* of size $O(\sqrt{n})$. In other words, for any planar graph $\mathcal{G} = (V, E)$ there exists a subset $S \subset V$ of size $O(\sqrt{n})$ with the following property: $V \setminus S$ can be split into subsets A and B with $|A| \leqslant 2n/3$ and $|B| \leqslant 2n/3$ such that there are no edges between A and B. This fundamental result was first proved in 1979 by Lipton and Tarjan [9] and has been refined in several ways, see e.g. [5, 6]. It has proved to be extremely useful for obtaining efficient divide-and-conquer algorithms for a large variety of problems on planar graphs.

In this paper we are interested in *geometric intersection graphs* in the plane. These are graphs whose node set corresponds to a set $\mathcal{D}$ of objects in the plane and that have an edge between two nodes iff the corresponding objects intersect. We will denote this intersection graph by $\mathcal{G}^{\times}(\mathcal{D})$. *(Unit) disk graphs* and *string graphs*—where the set $\mathcal{D}$ consists of (unit) disks and curves, respectively—are among the most popular types of intersection graphs. Unit disk graphs in particular have been studied extensively, because they serve as a model for wireless communication networks. It is well known that disk graphs are generalizations of planar graphs, because by the Circle Packing Theorem (also known as the Koebe–Andreev–Thurston Theorem) every planar graph is the intersection graph of a set of disks with disjoint interiors [10]. It is therefore natural to try to extend the Planar Separator Theorem to intersection graphs. A direct generalization is clearly impossible, however, since intersection graphs can contain arbitrarily large cliques. There are ways to still obtain separator theorems for intersection graphs. One can, for example, allow the size of the separator to depend on m, the number of edges, instead of on the number of vertices. It's known that any string graph admits a separator of size $O(\sqrt{m})$ [7].

Recently, De Berg *et al.* [3] introduced *clique-based separators*. A clique-based separator is a (balanced) separator consisting of cliques, and its size is not measured as the total number of nodes of the cliques but as the total number of cliques it contains. De Berg *et al.* showed that the intersection graph of a set $\mathcal{D}$ of convex fat objects in the plane admits a

40th European Workshop on Computational Geometry, Ioannina, Greece, March 13–15, 2024.
This is an extended abstract of a presentation given at EuroCG'24. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

separator consisting of $O(\sqrt{n})$ cliques. Clique-based separators are useful because cliques can be handled efficiently for many problems. De Berg *et al.* used their separators to solve many classic graph problems, including Independent Set, Dominating Set and more.

De Berg *et al.* [4] proved clique-based separator theorems for various other classes of objects, including map graphs and intersection graphs of pseudo-disks. They also showed that intersection graphs of *geodesic disks inside a simple polygon*—that is, geodesic disks induced by the standard shortest-path metric inside the polygon—admit a clique-based separator consisting of $O(n^{2/3})$ cliques. They left the case of geodesic disks in a polygon with holes as an open problem. We note that string graphs, which subsume the class of intersection graphs of geodesic disks, do *not* admit clique-based separators of sublinear size, since string graphs can contain arbitrarily large bipartite cliques.

Our results. We show that intersection graphs of geodesic disks in a polygon with holes admit a clique-based separator consisting of a sublinear number of cliques. Our result is actually much more general, as it shows that *for any well-behaved shortest-path metric d* defined on a path-connected and closed subset $F \subset \mathbb{R}^2$, a set $\mathcal{D}$ of geodesic disks with respect to that metric admits a separator consisting of $O(n^{4/5})$ cliques. Roughly speaking, we call a metric well-behaved if any two shortest paths under that metric meet at finitely many connected components, and any two disjoint components have some minimum positive clearance between them. This includes the shortest-path metric defined by a set of (possibly curved) obstacles in the plane, the shortest-path metric defined on a terrain, and the shortest-path metric among weighted regions in the plane. (The formal requirements on a well-behaved shortest-path metric can be found in the full version of the paper.) Note that we do not require shortest paths to be unique, nor do we put a bound on the number of intersections between the boundaries of two geodesic disks, nor do we require the metric space to have bounded doubling dimension.

The idea of our new approach is as follows. Recall that the *ply* of $\mathcal{D}$ is defined as the maximum number of objects from $\mathcal{D}$ with a common intersection. We first reduce the ply of the set $\mathcal{D}$ by removing all cliques of size $\Omega(n^{1/5})$, thus obtaining a set $\mathcal{D}^*$ of ply $O(n^{1/5})$. (The removed cliques will eventually be added to the separator. A similar preprocessing step was used in [4] to handle pseudo-disks.) The remaining arrangement can still be arbitrarily complex, however. To overcome this, we ignore the arrangement induced by the disks, and instead focus on the realization of the graph $\mathcal{G}^\times(\mathcal{D}^*)$ obtained by drawing a shortest path π_{ij} between the centers of any two intersecting disks $D_i, D_j \in \mathcal{D}^*$. We then prove that the number of edges of $\mathcal{G}^\times(\mathcal{D}^*)$ must be $O(n^{8/5})$; otherwise there will be an intersection point of two shortest paths $\pi_{ij}, \pi_{k\ell}$ that has large ply, which is not possible due to the preprocessing step. Since the number of edges of $\mathcal{G}^\times(\mathcal{D}^*)$ is $O(n^{8/5})$ we can use the separator result on string graphs to obtain a separator of size $O(\sqrt{n^{8/5}}) = O(n^{4/5})$. Adding the cliques that were removed in the preprocessing step then yields a separator consisting of $O(n^{4/5})$ cliques. In the full version we show how to improve the bound to $O(n^{3/4+\varepsilon})$ using a bootstrapping scheme. In this EuroCG paper, however, we only prove the weaker bound of $O(n^{4/5})$.

Clique-based separators give sub-exponential algorithms for Maximum Independent Set, Feedback Vertex Set, and q-Coloring for constant q [4]. When using our clique-based separator, the running times for Maximum Independent Set and Feedback Vertex Set are inferior to what is known for string graphs. For q-Coloring with $q \geqslant 4$ this is not the case, since it does not admit a sub-exponential algorithm, assuming ETH [2]. Our clique-based separator, however, yields an algorithm for geodesic disks running in $2^{O(n^{4/5})}$ time, if the boundaries of the disks D_i can be computed in polynomial time. In the full version we also show how to obtain an efficient distance-oracle for intersection graphs of geodesic disks.

2 A clique-based separator for geodesic disks in $\mathbb{R}^2$

Let d be a metric defined on a closed path-connected subset $F \subset \mathbb{R}^2$, and let $\mathcal{D} = \{D_1, \ldots, D_n\}$ be a set of geodesic disks in F, with respect to the metric d. Thus each disk D_i is defined as $D_i := \{q \in F : d(q, p_i) \leqslant r_i\}$, where $p_i \in F$ is the center of D_i an $r_i \geqslant 0$ is its radius. Let $\mathcal{D}_0 := \mathcal{D}$ and let $\mathcal{G}^{\times}(\mathcal{D}_0)$ denote the intersection graph of $\mathcal{D}_0$. We denote the set of edges of $\mathcal{G}^{\times}(\mathcal{D}_0)$ by E and define $m := |E|$.

2.1 Constructing the separator

We proceed in three steps: in a preprocessing step we reduce the ply of the set of disks we work with, in the second step we prove that if the ply is sublinear then the number of edges in the intersection graph is subquadratic, and in the third step we construct the separator.

Step 1: Reducing the ply. For a point $p \in F$, let $\mathcal{D}_0(p) := \{D_i \in \mathcal{D}_0 : p \in D_i\}$ be the set of disks from $\mathcal{D}_0$ containing p—note that $\mathcal{D}_0(p)$ forms a clique in $\mathcal{G}^{\times}(\mathcal{D}_0)$—and define $\text{ply}(p) := |\mathcal{D}_0(p)|$ to be the ply of p with respect to $\mathcal{D}_0$. The ply of the set $\mathcal{D}_0$ is defined as $\text{ply}(\mathcal{D}_0) := \max\{p \in F : \text{ply}(p)\}$.

We reduce the ply of $\mathcal{D}_0$ in the following greedy manner. Let α be a fixed constant with $0 < \alpha < 1$. In the basic construction we will use $\alpha = 1/5$, but in our bootstrapping scheme we will work with other values as well. We check whether there exists a point p such that $|\mathcal{D}_0(p)| \geqslant \frac{1}{4}n^{\alpha}$. If so, we remove $\mathcal{D}_0(p)$ from $\mathcal{D}_0$ and add it as a clique to $\mathcal{S}$. We repeat this process until $\text{ply}(\mathcal{D}_0) < \frac{1}{4}n^{\alpha}$. Thus in the first step at most $4n^{1-\alpha}$ cliques are added to $\mathcal{S}$.

To avoid confusion with our initial set $\mathcal{D}_0$, we denote the set of disks remaining at the end of Step 1 by $\mathcal{D}_1$ and we denote the set of edges of $\mathcal{G}^{\times}(\mathcal{D}_1)$ by E_1.

Step 2: Bounding the size of E_1. To bound the size of E_1, we draw a shortest path π_{ij} between the centers p_i, p_j of every two intersecting disks $D_i, D_j \in \mathcal{D}_1$, thus obtaining a geometric realization of the graph $\mathcal{G}^{\times}(\mathcal{D}_1)$. Slightly abusing notation, we will not distinguish between an edge (D_i, D_j) in $\mathcal{G}^{\times}(\mathcal{D}_1)$ and its geometric realization π_{ij}. Let $\Pi(\mathcal{D}_1) := \{\pi_{ij} : (D_i, D_j) \in E_1\}$ be the resulting set of paths. To focus on the main idea behind our proof we will assume that $\Pi(\mathcal{D}_1)$ is a *proper path set*: a set of paths such that any two paths $\pi_{ij}, \pi_{k\ell}$ have at most two points in common, each intersection point is either a shared endpoint or a proper crossing, and no proper crossing coincides with another proper crossing or with an endpoint. The proof also works without this assumption, given that the metric we are working with is well-behaved. We defer this discussion to the full version. We need the following result about the number of crossings in dense graphs, known as the Crossing Lemma [1, 8]. The term *planar drawing* here refers to a drawing where no edge interior passes through a vertex and all intersections are proper crossings. Thus it applies to a proper path set.

▶ **Lemma 2.1** (Crossing Lemma). *There exists a constant $c > 0$, such that every planar drawing of a graph with n vertices and $m \geqslant 4n$ edges contains at least $c\frac{m^3}{n^2}$ crossings.*

Using the Crossing Lemma we will show that $|E_1| = O(n^{\frac{3+\alpha}{2}})$, as follows. If $|E_1| = O(n^{\frac{3+\alpha}{2}})$ does not hold, then by the Crossing Lemma there must be many crossings between the edges $\pi_{ij} \in E_1$. We will show show that this implies that there is a crossing of ply greater than $\frac{1}{4}n^{\alpha}$, thus contradicting that $\text{ply}(\mathcal{D}_1) < \frac{1}{4}n^{\alpha}$. We now make this idea precise.

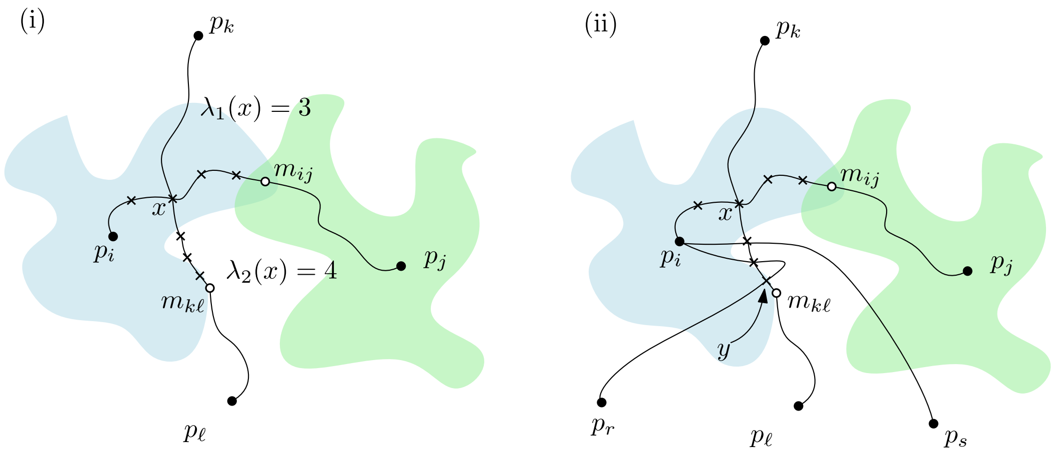

Figure 1 (i) A labeling of a crossing $x \in \mathcal{X}$. (ii) The crossing y is assigned to D_i four times.

▶ **Lemma 2.2.** *Let $\mathcal{G}^\times(\mathcal{D}_1) = (\mathcal{D}_1, E_1)$ be the intersection graph of a set $\mathcal{D}_1$ of disks such that $\text{ply}(\mathcal{D}_1) < \frac{1}{4}n^\alpha$. Then $|E_1| \leqslant \sqrt{\frac{4}{c}} \cdot n^{\frac{3+\alpha}{2}}$, where c is the constant appearing in the Crossing Lemma.*

Proof. Consider the proper path set $\Pi(\mathcal{D}_1)$ and let $\mathcal{X}$ be the set of crossings between the paths in $\Pi(\mathcal{D}_1)$. Assume for a contradiction that $|E_1| > \sqrt{\frac{4}{c}} \cdot n^{\frac{3+\alpha}{2}}$. We will show that then there has to exist a crossing $x \in \mathcal{X}$ of ply at least $\frac{n^\alpha}{4}$, which contradicts that $\text{ply}(\mathcal{D}_1) < \frac{1}{4}n^\alpha$.

We start by giving a lower bound on the total ply of all crossings in the drawing. To this end, we split each edge $\pi_{ij} \in E_1$ in two *half-edges* as follows. For two points $x, y \in \pi_{ij}$, let $\pi_{ij}[x, y]$ denote the subpath of π_{ij} between x and y. Recall that p_i is the center of disk D_i. We now pick an arbitrary point $m_{ij} \in \pi_{ij} \cap (D_i \cap D_j)$ and split π_{ij} at m_{ij} into a half-edge $\pi_{ij}[p_i, m_{ij}]$ connecting p_i to m_{ij} and a half-edge $\pi_{ij}[p_j, m_{ij}]$ connecting p_j to m_{ij}. For brevity, we will denote these two half-edges by h_{ij} and h_{ji}, respectively. Clearly each half-edge has length at most the radius of the disk it lies in, and so $h_{ij} \subset D_i$ and $h_{ji} \subset D_j$. We denote the resulting set of half-edges by $\mathcal{E}_1$.

We label each crossing $x \in \mathcal{X}$ with an unordered pair of integers $\{\lambda_1(x), \lambda_2(x)\}$, defined as follows: if x is the crossing between the half-edges $h_{ij}, h_{k\ell}$, then $\lambda_1(x)$ is the number of crossings contained in $\pi_{ij}[x, m_{ij}]$ and $\lambda_2(x)$ is the number of crossings contained in $\pi_{k\ell}[x, m_{k\ell}]$; see Fig. 1(i). This labeling is useful to obtain a rough bound on the total ply of all crossings, because of the following observation, which immediately follows from the triangle inequality.

Observation 1. Consider a crossing $x = h_{ij} \cap h_{k\ell}$. If $d(x, m_{k\ell}) \leqslant d(x, m_{ij})$ then all crossings $y \in \pi_{k\ell}[x, m_{k\ell}]$ are contained in D_i, and otherwise all crossings $y \in \pi_{ij}[x, m_{ij}]$ are contained in D_k.

Let $K := \sum_{x\in\mathcal{X}} \text{ply}(x)$ denote the total ply of all crossings. The following claim bounds K in terms of the labels $\{\lambda_1(x), \lambda_2(x)\}$.

Claim 1. $K > \frac{1}{2|\mathcal{D}_1|} \sum_{x\in\mathcal{X}} \min\{\lambda_1(x), \lambda_2(x)\}$.

Proof. Define $K(D_i) := \big|\{x \in \mathcal{X} : x \in D_i\}\big|$ to be the contribution of D_i to the total ply K, and note that

$$K = \sum_{x\in\mathcal{X}} \mathrm{ply}(x) = \sum_{x\in\mathcal{X}} \left|\{D_i \in \mathcal{D}_1 : x \in D_i\}\right| = \sum_{D_i\in\mathcal{D}_1} \left|\{x \in \mathcal{X} : x \in D_i\}\right| = \sum_{D_i\in\mathcal{D}_1} K(D_i).$$

Now consider a crossing $x = h_{ij} \cap h_{k\ell}$. If $d(x, m_{k\ell}) \leqslant d(x, m_{ij})$ then we assign x to D_i, and otherwise we assign x to D_k. Let $\mathcal{X}(D_i)$ be the set of crossings assigned to D_i. By Observation 1 and the definition of the label $\{\lambda_1(x), \lambda_2(x)\}$, the disk D_i contains at least $\min\{\lambda_1(x), \lambda_2(x)\}$ crossings $y \in \pi_{k\ell}[x, m_{k\ell}]$. Thus, summing over all crossings $x \in \mathcal{X}(D_i) \cap h_{ij}$ and all half-edges h_{ij} incident to D_i, we find that

$$K(D_i) \geqslant \frac{1}{2\deg(D_i)} \sum_{h_{ij}} \sum_{x\in\mathcal{X}(D_i)\cap h_{ij}} \min\{\lambda_1(x), \lambda_2(x)\} > \frac{1}{2|\mathcal{D}_1|} \sum_{x\in\mathcal{X}(D_i)} \min\{\lambda_1(x), \lambda_2(x)\},$$

where $\deg(D_i)$ denotes the degree of D_i in $\mathcal{G}^\times(\mathcal{D}_1)$. The factor $\frac{1}{2\deg(D_i)}$ arises because a crossing $y \in h_{k\ell}$ can be counted up to $2\deg(D_i)$ times in the expression $\sum_{x\in\mathcal{X}(D_i)} \min\{\lambda_1(x), \lambda_2(x)\}$, namely at most twice for every half-edge incident to D_i; see Fig. 1(ii). (Twice, because a pair of paths in a proper path set may cross twice.) Since each crossing is assigned to exactly one set $\mathcal{X}(D_i)$, we obtain

$$K = \sum_{D_i\in\mathcal{D}_1} K(D_i) > \sum_{D_i\in\mathcal{D}} \left(\frac{1}{2|\mathcal{D}_1|} \sum_{x\in\mathcal{X}(D_i)} \min\{\lambda_1(x), \lambda_2(x)\} \right) = \frac{1}{2|\mathcal{D}_1|} \sum_{x\in\mathcal{X}} \min\{\lambda_1(x), \lambda_2(x)\}.$$

$\triangleleft$

From the Crossing Lemma and our initial assumption that $|E_1| > \sqrt{\frac{4}{c}} \cdot n^{\frac{3+\alpha}{2}}$, we have that

$$|\mathcal{X}| \geqslant c\frac{|E_1|^3}{n^2} > 4|E_1|n^{1+\alpha} = 2|\mathcal{E}_1|n^{1+\alpha}. \tag{1}$$

In order to get a rough bound for $\sum_{x\in\mathcal{X}} \min\{\lambda_1(x), \lambda_2(x)\}$, we will ignore crossings with small labels, while ensuring that we don't ignore too many crossings in total. More precisely, for every half-edge h_{ij}, we disregard its first $n^{1+\alpha}$ crossings, starting from the one closest to m_{ij}. We let $\mathcal{X}^*$ denote the set of remaining crossings. Note that $|\mathcal{X}^*| \geqslant |\mathcal{X}| - |\mathcal{E}_1|n^{1+\alpha}$, and $\min\{\lambda_1(x), \lambda_2(x)\} \geqslant n^{1+\alpha}$ for every $x \in \mathcal{X}^*$. Therefore

$$K > \frac{1}{2|\mathcal{D}_1|} \sum_{x\in\mathcal{X}^*} \min\{\lambda_1(x), \lambda_2(x)\} = \frac{1}{2|\mathcal{D}_1|} \cdot |\mathcal{X}^*| \cdot n^{1+\alpha} \geqslant \frac{\left(|\mathcal{X}| - |\mathcal{E}_1|n^{1+\alpha}\right) n^{1+\alpha}}{2|\mathcal{D}_1|} \geqslant \frac{|\mathcal{X}|}{4} n^{\alpha}.$$

This means that there exists a crossing $x \in \mathcal{X}$ of ply at least $\frac{1}{4}n^\alpha$, which contradicts the condition of the lemma and thus finishes the proof. ◀

Step 3: Applying the separator theorem for string graphs. Lee's separator theorem for string graphs [7] states that any string graph on m edges admits a balanced separator of size $O(\sqrt{m})$. It is well known that any intersection graph of path-connected sets in the plane is a string graph. Hence, we can apply Lee's result to $\mathcal{G}^\times(\mathcal{D}_1)$ which has $O(n^{\frac{3+\alpha}{2}})$ edges. Thus, $\mathcal{G}^\times(\mathcal{D}_1)$ has a separator of size $O(n^{\frac{3+\alpha}{4}})$. By adding the vertices of this separator as singletons to our separator $\mathcal{S}$ then, together with the cliques added in Step 1, we obtain a separator consisting of $O(n^{\frac{3+\alpha}{4}} + n^{1-\alpha})$ cliques. Picking $\alpha = 1/5$, and anticipating the extension to the case where $\Pi(\mathcal{D}_1)$ is not a proper path set, we obtain the following result.

▶ **Theorem 2.3.** *Let d be a well-behaved shortest-path metric on a closed and path-connected subset $F \subset \mathbb{R}^2$ and let $\mathcal{D}$ be a set of n geodesic disks with respect to the metric d. Then $\mathcal{G}^{\times}(\mathcal{D})$ has a balanced clique-based separator consisting of $O(n^{4/5})$ cliques.*

De Berg *et al.* [4] showed that if a class of graphs admits a clique-based separator consisting of $S(n)$ cliques, and the separator can be constructed in polynomial time, then one can solve q-Coloring in $2^{O(S(n))}$ time. Note that if we can compute the boundaries ∂D_i in polynomial time,[1] then we can also compute our separator in polynomial time. Indeed, we can then compute $\mathcal{G}^{\times}(\mathcal{D})$ and the arrangement of the disk boundaries in polynomial time, which allows us to do Step 1 (reducing the ply) in polynomial time. Since a separator for string graphs can be computed in polynomial time [7], this is easily seen to imply that the separator construction runs in polynomial time. Thus we obtain the following result.

▶ **Corollary 2.4.** *Let d be a shortest-path metric on a connected subset $F \subset \mathbb{R}^2$ and let $\mathcal{D}$ be a set of n geodesic disks with respect to the metric d, where d is such that the boundaries ∂D_i of the disks in $\mathcal{D}$ can be computed in polynomial time. Let $q \geqslant 1$ and $\varepsilon > 0$ be fixed constants. Then q-Coloring can be solved in $2^{O(n^{\frac{4}{5}})}$ time on $\mathcal{G}^{\times}(\mathcal{D})$.*

References

1. M. Ajtai, V. Chvátal, M.M. Newborn, and E. Szemerédi. Crossing-free subgraphs. In Peter L. Hammer, Alexander Rosa, Gert Sabidussi, and Jean Turgeon, editors, *Theory and Practice of Combinatorics*, volume 60 of *North-Holland Mathematics Studies*, pages 9–12. North-Holland, 1982. doi:doi.org/10.1016/S0304-0208(08)73484-4.
2. 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.
3. 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.
4. Mark de Berg, Sándor Kisfaludi-Bak, Morteza Monemizadeh, and Leonidas Theocharous. Clique-based separators for geometric intersection graphs. *Algorithmica*, 85(6):1652–1678, 2023. doi:10.1007/S00453-022-01041-8.
5. Hristo Djidjev and Shankar M. Venkatesan. Reduced constants for simple cycle graph separation. *Acta Informatica*, 34(3):231–243, 1997. doi:10.1007/s002360050082.
6. Jacob Fox and János Pach. Separator theorems and Turán-type results for planar intersection graphs. *Advances in Mathematics*, 219(3):1070–1080, 2008. doi:doi.org/10.1016/j.aim.2008.06.002.
7. 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.
8. Thomas Leighton. *Complexity Issues in VLSI.* Foundations of Computing Series. MIT Press, 2003.
9. 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.
10. William Thurston. *The geometry and topology of 3-manifolds.* Princeton Lecture Notes, 1978–1981.

[1] Typically the time to do this would not only depend on n, but also on the complexity of F and the distance function d. For simplicity we state our results in terms of n only. (This, of course, puts restrictions on the complexity of F and d.)