

Sweeping Arrangements of Non-Piercing Regions in the Plane

Suryendu Dalal ✉

IIT-Delhi, India

Rahul Gangopadhyay ✉

Moscow Institute of Physics and Technology, Russia

Rajiv Raman ✉

IIT-Delhi, India

Saurabh Ray ✉

NYU Abu Dhabi, UAE

Abstract

Let Γ be a finite set of Jordan curves in the plane. For any curve $\gamma \in \Gamma$, we denote the bounded region enclosed by γ as $\tilde{\gamma}$. We say that Γ is a non-piercing family if for any two curves $\alpha, \beta \in \Gamma$, $\tilde{\alpha} \setminus \tilde{\beta}$ is a connected region. A non-piercing family of curves generalizes a family of 2-intersecting curves in which each pair of curves intersect in at most two points. Snoeyink and Hershberger (“Sweeping Arrangements of Curves”, SoCG ’89) proved that if we are given a family Γ of 2-intersecting curves and a *sweep* curve $\gamma \in \Gamma$, then the arrangement can be *swept* by γ while always maintaining the 2-intersecting property of the curves. We generalize the result of Snoeyink and Hershberger to the setting of non-piercing curves. We show that given an arrangement of non-piercing curves Γ , and a sweep curve $\gamma \in \Gamma$, the arrangement can be swept by γ so that the arrangement remains non-piercing throughout the process. We also give a shorter and simpler proof of the result of Snoeyink and Hershberger, and give an eclectic set of applications.

2012 ACM Subject Classification Theory of computation \rightarrow Computational geometry

Keywords and phrases Sweeping, Pseudodisks, Discrete Geometry, Topology

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.45

Related Version *Full Version*: <https://arxiv.org/abs/2403.16474>

Funding *Rahul Gangopadhyay*: Rahul Gangopadhyay was supported by a grant for research centers in the field of artificial intelligence, provided by the Analytical Center for the Government of the Russian Federation in accordance with the subsidy agreement (agreement identifier 000000D730321P5Q0002) and the agreement with the Moscow Institute of Physics and Technology dated November 1, 2021 No. 70-2021-00138.

Acknowledgements We would like to thank the anonymous reviewers for their constructive comments that improved the presentation of the paper.

1 Introduction

A fundamental algorithmic design technique in Computational Geometry introduced by Bentley and Ottmann [4] is *line sweep*. In the original context the technique was invented, namely reporting all intersections among n line segments in the plane. The basic idea is to move a vertical line from $x = -\infty$ to $x = +\infty$, stopping at *event points* (like end-points of segments or intersections of two segments) where updates are made. Since then, the line sweep technique has found a wide variety of applications in Computational Geometry including polygon triangulation, computing Voronoi diagrams and its dual Delaunay triangulations, and computing the volume of a union of regions to name a few. The method can also



© Suryendu Dalal, Rahul Gangopadhyay, Rajiv Raman, and Saurabh Ray;
licensed under Creative Commons License CC-BY 4.0


40th International Symposium on Computational Geometry (SoCG 2024).

Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 45; pp. 45:1–45:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



be generalized to higher dimensions. See the classic books on Computational Geometry [10, 17, 11] for more applications. Besides applications in Computational Geometry, it is also a useful tool in combinatorial proofs. For a simple example, a one dimensional sweep can be used to show that the clique cover number and maximum independent set size in interval graphs are equal. Several other examples appear in Section 3 of this paper. In most applications, the sweep line moves continuously and covers the plane so that each point in the plane is on the sweep line exactly once. There are also applications where instead of sweeping with a line, we sweep with a closed curve like a circle. For instance we can start with an infinite radius circle and shrink it continuously to its center.

Edelsbrunner and Guibas [12] introduced the *topological sweep* technique and showed the advantage of sweeping with a *topological line*, i.e., curve that at any point intersects each segment at most once rather than a rigid straight line. They obtained better algorithms to report all intersection points between lines, which by duality improved results known at that time for many problems on point configurations. Chazelle and Edelsbrunner [8] were also able to adapt the topological sweep “in 20 easy pieces” to report the k intersections of n given lines segments in $O(n \log n + k)$ time.

Thus, we can generalize the sweeping technique to sweep a set of *pseudolines*¹ with a pseudoline. We can generalize further, and do topological sweep with a closed curve. Generalizing from a set of circles in the plane, where any pair of circles intersect in at most two points, a collection of simple closed curves is called an *arrangement of pseudocircles* iff any pair of curves intersect in at most two points. Note that for any two disks defined by circles or pseudocircles, D_1, D_2 , the difference $D_1 \setminus D_2$, is a path-connected set.

Showing that it is possible to sweep a collection of pseudolines with one of the pseudolines, and a set of pseudocircles with one of the pseudocircles is much more challenging than sweeping with a line or circle. Snoeyink and Hershberger [21] showed in a celebrated paper that we can sweep an arrangement of 2-intersecting curves with any curve in the family. The result of Snoeyink and Hershberger has found several applications both in Computational Geometry as well as in combinatorial proofs. See for example, [13, 2, 5, 1, 9, 3, 20] and references therein.

It is natural to attempt to generalize the result of Snoeyink and Hershberger to k -intersecting curves for $k > 2$, i.e., sweep a set of k -intersecting curves with a sweep curve that maintains the invariant that the arrangement is k -intersecting throughout the sweep. Unfortunately, Snoeyink and Hershberger [21] show that this is not possible for $k > 3$, and ask at the end of their paper what intersection property the sweep satisfies if we sweep an arrangement of k -intersecting curves. While we do not fully answer their question, we extend their result to show that we can maintain a more general property than 2-intersection while sweeping. Specifically, we show that if the input curves satisfy a topological condition of being *non-piercing*, then we can sweep with a curve so that the arrangement satisfies this property throughout the sweep. For a Jordan curve γ , let $\tilde{\gamma}$ denote the bounded region defined by γ . Two Jordan curves α, β are said to be *non-piercing* if $\tilde{\alpha} \setminus \tilde{\beta}$ is a path connected region, and a family of curves is non-piercing if they are non-piercing pairwise. The main theorem we prove in this paper is the following.

► **Theorem 1** (Sweeping non-piercing arrangements). *Let Γ be a finite non-piercing family of curves. Given any $\gamma \in \Gamma$ and a point $P \in \tilde{\gamma}$, we can sweep Γ with γ so that at any point of time during the sweep, the curves remain non-piercing and P remains within $\tilde{\gamma}$.*

¹ A set \mathcal{L} of bi-infinite curves in the plane is a collection of pseudolines the curves in \mathcal{L} pairwise intersect at most once.

We also give an alternative, and somewhat simpler proof of the fundamental result of Snoeyink and Hershberger. While our result is not significantly simpler, we believe it leads to a cleaner analysis using two conceptual tools, namely *minimal lens bypassing* and *minimal triangle bypassing* that may help in other settings involving pseudodisks or non-piercing regions. Further, the proof of Theorem 1 is quite simple compared to the result of Snoeyink and Hershberger, and even our simpler proof of their result (Theorem 14). Hence, the main difficulty in sweeping seems to be maintaining the 2-intersection property.

For many algorithmic applications, especially for several packing and covering problems, the restriction that an arrangement is non-piercing is not harder than the corresponding problems for pseudodisks. For example, hypergraphs defined by points and non-piercing regions enjoy a linear *shallow-cell complexity* [18] and therefore, admit ϵ -nets of linear size and $O(1)$ -approximation algorithms for covering problems via *quasi-uniform* sampling [22, 7]. In particular, it is plausible that the non-piercing condition is the most elementary condition required for which results that work for pseudodisks extend. For example Sariel Har-Peled [14], and Chan and Grant [6] showed that for several covering problems involving simple geometric regions that are not *non-piercing*, several problems that admit a PTAS for non-piercing regions [18] become APX-hard.

2 Preliminaries

Let \tilde{S}^1 denote a circle centered at the origin and oriented counter-clockwise around the origin. An oriented Jordan curve γ is a continuous injective map from $\tilde{S}^1 \rightarrow \mathbb{R}^2$ that respects the orientation of \tilde{S}^1 . i.e., the origin is mapped to a point in the region bounded by γ . Unless otherwise stated, by a *curve*, we will mean an oriented Jordan curve.

In the following, we represent by Γ a finite set of oriented Jordan curves. For any $\gamma \in \Gamma$, we denote by $\tilde{\gamma}$, the bounded region defined by γ . Observe that $\tilde{\gamma}$ is simply-connected².

A *traversal* of a curve is a walk starting at an arbitrary point on the curve, walking in the direction of its orientation and returning to the starting point. Let $\mathcal{A}(\Gamma)$ denote the arrangement of the curves in Γ . We assume throughout the paper that the curves are in general position, i.e., no three curves intersect at a point, any pair of curves intersect in a finite number of points, and they intersect transversally (i.e., they cross) at these points. The arrangement consists of vertices (0-dimensional faces), where two curves intersect, arcs (1-dimensional faces) which are segments of curves between vertices and cells (2-dimensional faces) which are the connected regions of $\mathbb{R}^2 \setminus \Gamma$. The boundary of a cell thus consists of vertices and arcs. Each arc is a portion of a curve, and each vertex is the intersection of two of the curves. The set of curves which contribute one or more arcs to the boundary of a cell are said to define the cell. If a curve α contributes an arc to the boundary of a cell C , we say that C *lies on* α .

A *digon* is a cell D in the arrangement of any two curves $\alpha, \beta \in \Gamma$ so that both α and β contribute one arc to the boundary of D . Note that D is not necessarily a cell in the arrangement of all curves in Γ . If a cell in $\mathcal{A}(\Gamma)$ is a digon, we call it a *digon cell*. A digon defined by curves α and β is called a *lens* if it is contained in both $\tilde{\alpha}$ and $\tilde{\beta}$ and it is called a *negative lens* if it is not contained in either $\tilde{\alpha}$ or $\tilde{\beta}$. A *triangle* is a cell T in the arrangement of three curves in Γ so that each curve contributes one arc to the boundary of T . A triangle need not be a cell in $\mathcal{A}(\Gamma)$. If a cell in $\mathcal{A}(\Gamma)$ is a triangle, we call it a *triangle cell*.

² A region $r \subseteq \mathbb{R}^2$ is said to be *simply-connected* if any closed loop can be continuously deformed to a point inside the region.

Given an arrangement of oriented Jordan curves Γ and a sweeping curve $\gamma \in \Gamma$, a property Π satisfied by the curves in Γ , and a point $P \in \tilde{\gamma}$, a sweep is a continuous movement of γ until the region it bounds shrinks to the point P so that (i) for any two times $t > t'$, the curve γ at time t is contained in the bounded region of γ at time t' , (ii) at each point in time the arrangement of curves satisfies the property Π . In particular, if the property Π is that the curves are non-piercing, we want to ensure that the curves remain non-piercing throughout the process.

Pseudocircles and Non-piercing regions. If the curves in Γ intersect pairwise at at most two points, we call Γ a family of *pseudocircles* and in this case the set $\{\tilde{\gamma} : \gamma \in \Gamma\}$ is called a family of *pseudodisks*. Let α and β be two curves. For $k \in \mathbb{N}$, we say that α and β are *k-intersecting* if they intersect in at most k points. We say that α and β are *non-piercing* if $\tilde{\alpha} \setminus \tilde{\beta}$, and $\tilde{\beta} \setminus \tilde{\alpha}$ are path-connected. Otherwise, they are said to be *piercing*. A finite collection of curves Γ is said to be *non-piercing* if the curves in Γ are pairwise non-piercing.

3 Applications

In this section, we describe several applications of Theorem 1. We omit the proofs of some of the theorems in this section since they only require replacing the result of Snoeyink and Hershberger (Theorem 14) with our result (Theorem 1).

Every arrangement of non-piercing regions contains a “small” non-piercing region. Pinchasi [16] proved that an arrangement of pseudodisks Γ contains a pseudodisk γ s.t. γ is intersected by at most 156 disjoint pseudodisks in $\Gamma \setminus \{\gamma\}$. Using Theorem 1, Pinchasi’s result extends immediately to non-piercing regions.

► **Theorem 2.** *Let Γ be an arrangement of non-piercing curves in the plane. Then, there is a curve $\gamma \in \Gamma$ s.t. the number of curves in Γ whose regions are disjoint and intersect $\tilde{\gamma}$ is at most 156.*

Multi-hitting set with non-piercing regions. Raman and Ray [19] studied Hitting set and Set cover problems with demands for set systems defined by a set of points and set of non-piercing regions in the plane. Their proofs rely on showing the existence of suitable graphs. For one of the results they used the result of Snoeyink and Hershberger and consequently it applied only to pseudodisks. Using Theorem 1, it can be extended to non-piercing regions.

► **Theorem 3.** *Given an instance (P, \mathcal{S}) where P is a set of points in the plane and \mathcal{S} is a collection of non-piercing regions in the plane with demands $d : \mathcal{S} \rightarrow \mathbb{N}$, there is a PTAS for the Multi-hitting set problem, i.e., selecting the smallest size subset $Q \subseteq P$ s.t. $\forall S \in \mathcal{S}, Q \cap S \geq d(S)$ when the demands are bounded above by a constant, and otherwise there is a $(2 + \epsilon)$ -approximation algorithm.*

Number of hyperedges defined by lines and non-piercing regions. Keller et al., [15] considered the hypergraph $(\mathcal{L}, \mathcal{D})$, where the elements of the hypergraph are defined by a set \mathcal{L} of lines in the plane, and each pseudodisk $D \in \mathcal{D}$ defines a hyperedge consisting of the lines in \mathcal{L} intersecting a pseudodisk. The authors prove that the number of hyperedges in such a hypergraph is $O(|\mathcal{L}|^3)$. By using Theorem 1 in place of the result of Snoeyink and Hershberger (Theorem 14), and following their result, we obtain the same bound for the hypergraph defined by lines and non-piercing regions.

► **Theorem 4.** Let $(\mathcal{L}, \mathcal{D})$ be a hypergraph defined by a set \mathcal{L} , where each hyperedge is defined by the set of lines intersecting a non-piercing region. Then, the number of hyperedges of size t in this hypergraph is $O(|\mathcal{L}|^2)$, and the total number of hyperedges is $O(|\mathcal{L}|^3)$. Both bounds are tight already for pseudodisks.

Construction of Supports. Raman and Ray[18] proved the following:

► **Theorem 5** (Support for non-piercing regions). Let \mathcal{H} be a family of non-piercing regions and let P be any set of points in the plane. There exists a planar graph $G = (P, E)$ s.t. for any $H \in \mathcal{H}$ the subgraph of G induced by $H \cap P$ is connected.

Using Theorem 1, we obtain a significantly simpler proof.

4 Sweeping

Let Γ be a set of curves in the plane and let $\gamma \in \Gamma$ be one of the curves called the *sweep curve*. By “sweeping of $\mathcal{A}(\Gamma)$ by γ ”, we mean the process of continuously shrinking $\tilde{\gamma}$ over a finite duration of time until there are no intersection points among the curves in Γ in the interior of $\tilde{\gamma}$. In general, we want to do this while maintaining some topological invariants during the sweep. Snoeyink and Hershberger [21] discretized this continuous sweeping process into a set of discrete operations each of which can be implemented as a continuous deformation of γ over a unit time interval. We follow their arguments and likewise describe a set of allowable discrete *sweeping operations*. Most of our operations are slight variations on their operations, except one new operation that we introduce. In order to describe the operations, we need one more definition. We say that a curve α is *visible* from the sweep curve γ if there is an arc $\tau(u, v)$ (called the *visibility arc*) joining a point u on γ with a point v on α whose interior lies in $\tilde{\gamma}$ and does not intersect any of the curves in Γ . We define the following discrete sweeping operations:

1. *Take a loop:* Let α be a curve that does not intersect γ and such that α is visible from γ via the visibility arc $\tau(u, v)$. We define “taking a loop” as the operation that modifies γ to γ' as follows: Taking two points u', u'' very close to u on either side to u on γ , we replace the segment of γ between u' and u'' by a curve between u' and u'' that lies very close to $\tau(u, v)$ and loops around v crossing α twice. See Figure 1. More formally, γ is modified to γ' so that $\tilde{\gamma}' = \tilde{\gamma} \setminus (\tau(u, v) \oplus B_\epsilon)$ where \oplus denotes Minkowski sum and B_ϵ is a ball of radius ϵ for an arbitrarily small ϵ . In the rest of the paper, we avoid such formal definitions and resort to figures for ease of exposition. However, our informal definitions using figures can easily be formalized.



■ **Figure 1** Taking a loop.

2. *Bypass a digon cell:* Suppose that a curve α and the sweep curve γ form digon cell D with vertices u and v . Then, we define the operation of “bypassing D ” as the modification of γ to a curve γ' that goes around the digon cell D as shown in Figure 2.



■ **Figure 2** Bypassing a digon cell.

3. *Bypass a triangle cell:* Let α and β be two curves that along with the sweep curve γ define a triangle cell T . We define the operation of “bypassing T ” as the modification of γ to a curve γ' which goes around T as show in Figure 3.



■ **Figure 3** Bypassing a triangle cell.

4. *Bypass a visible vertex:* Let v be a vertex lying inside $\tilde{\gamma}$ where the curves α and β intersect. If there is a visibility arc $\tau(u, v)$ from a point u on γ to v whose interior lies in the interior of $\tilde{\alpha}$ as well as $\tilde{\beta}$, then by definition, the interior of $\tau(u, v)$ also lies in the interior of $\tilde{\gamma}$. In this case, we refer to $\tau(u, v)$ as the *bypassability curve* for v . “Bypassing v ” is the modification of γ to a curve γ' that takes a loop around v as shown in Figure 4. Note that unlike the previous operations, the orientation of the curves are important here.



(a) Visibility curve $\tau(u, v)$ to v .

(b) The curve γ' after bypassing vertex v .

■ **Figure 4** Bypassing a visible vertex.


The following Lemma summarizes properties of the above operations that are intuitively obvious from the figures.

► **Lemma 6.** *Let Γ be an arrangement of non-piercing curves with sweep curve $\gamma \in \Gamma$. The operations (i) Take a loop, (ii) Bypass a digon cell, (iii) bypass a triangle cell, and (iv) bypass a visible vertex, leave the arrangement of curves non-piercing but can change the number of intersections between γ and other curves. Operation (i) increases the number of intersection points between γ and the curve α we take a loop on by 2. Operation (ii) decreases the number of intersections between γ and the other curve α defining the digon cell. Operation (iii) does not change the number of intersections between γ and other curves. Finally, operation (iv) increases by 2 the number of intersections between γ and each of the curves α, β defining the vertex being bypassed.*

► **Corollary 7.** *Let Γ be an arrangement of 2-intersecting curves with sweep curve $\gamma \in \Gamma$. Then, applying any of the three operations: taking a loop on a curve that does not intersect γ , bypassing a digon cell, or bypassing a triangle cell leaves the arrangement two-intersecting.*

5 Minimal Digons and Triangles

In this section, we describe a basic operation required in our proofs, namely that of bypassing *minimal digons* or *minimal triangles* in an arrangement $\mathcal{A}(\Gamma)$ of a set of non-piercing curves Γ . Recall that in general digons or triangles are not cells in the arrangement $\mathcal{A}(\Gamma)$. A digon L formed by curves α and β is minimal if it does not contain another digon. Equivalently, L is minimal if the intersection of any curve δ with L , consists of disjoint arcs each having one end-point on α and one end-point on β . Let L be a minimal digon formed by curves α and β with vertices u and v . The operation of bypassing L modifies the arcs of α and β as shown in Figure 5. Formally, the bypassing is done in two steps. First we replace the arc of α between u and v by the arc of β between u and v and vice versa so that u and v are now points of tangencies between α and β . Next we get rid of the tangencies by moving α and β slightly apart from each other around u and v . We also add *reference points* called u_α and u_β arbitrarily close to u so that α passes between u and u_α keeping u_β on the same side as u , and similarly β passes between u and u_β keeping u_α on the same side as u . We define the reference points v_α and v_β close to v analogously. Note that the reference points are not vertices in the modified arrangement. We claim that if a minimal digon in a non-piercing family is bypassed, the modified family obtained is still non-piercing.




■ **Figure 5** The operation of digon bypassing of curves α and β .

► **Lemma 8.** *Let L be a minimal digon formed by α and β in the arrangement $\mathcal{A}(\Gamma)$ as above. The modified family of curves obtained by bypassing L is non-piercing. Furthermore, the number of intersections between α and β decreases by 2 and for any other pair of curves the number of intersections does not change.*

The only new/changed cells in the modified arrangement obtained after bypassing L are those that contain one of the points $u, v, u_\alpha, u_\beta, v_\alpha$ or v_β . Note that there may also be cells that are identical to an old cell but their defining set of curves has changed - α is replaced by β or vice versa.

► **Lemma 9.** *Only the new cells containing one of the reference points $u_\alpha, u_\beta, v_\alpha$ or v_β may be newly created digon/triangle cells on the sweep curve γ in the arrangement obtained after bypassing L .*

Let T be a triangle on the sweep curve γ defined by the curves $\alpha, \beta \in \Gamma$ along with γ . We say that T is a “minimal triangle on γ ”, or “base γ ” if the intersection of any other curve $\delta \in \Gamma \setminus \{\alpha, \beta, \gamma\}$ with T consists of a set of disjoint arcs each with one end-point on α and one end-point on β . The operation of bypassing T modifies α and β so that they go around T and their intersection u on the boundary of T moves outside $\tilde{\gamma}$ as shown in Figure 6. We show next, that bypassing T leaves the arrangement non-piercing.




■ **Figure 6** The operation of minimal triangle bypassing of curves α and β .

► **Lemma 10.** *Let T be a minimal triangle in an arrangement Γ of non-piercing regions defined by curves α, β and γ . Then, bypassing T yields a non-piercing arrangement. The number of intersection points between any pair of curves does not change. However, the number of intersection points in $\tilde{\gamma}$ decreases by 1. Thus, the number of intersections in $\tilde{\gamma}$ decreases by 1.*

► **Corollary 11.** *Let T be a minimal triangle in an arrangement Γ of 2-intersecting defined by α, β and γ . Then, bypassing T yields a 2-intersecting arrangement.*

Let T be a minimal triangle on γ defined by α and β . Let u, v and w be the vertices of T defined by the pairs (γ, α) , (α, β) , and (β, γ) respectively. Note that u is a vertex of a unique lens defined by β and γ . Let us call the other vertex of this lens u' . We define w' similarly so that w and w' are vertices of a lens defined by α and γ .



■ **Figure 7** The operation of minimal triangle bypassing of curves α and β .

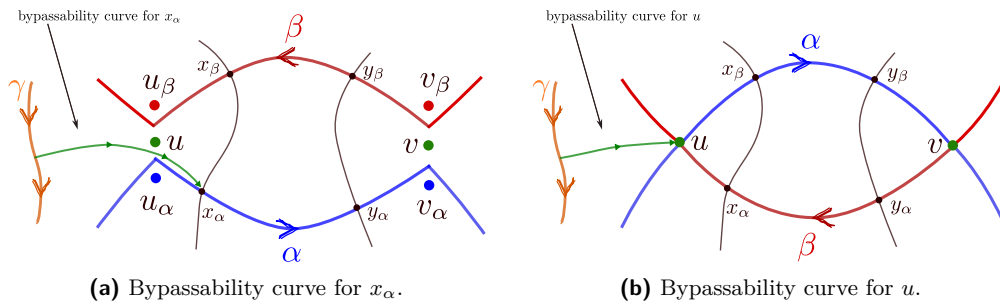
► **Lemma 12.** *If there is at least one other curve δ that intersects the minimal triangle T , then only the cells containing the reference points v_α or v_β may be digon/triangle cells.*

6 Sweeping Non-piercing regions

► **Theorem 1** (Sweeping non-piercing arrangements). *Let Γ be a finite non-piercing family of curves. Given any $\gamma \in \Gamma$ and a point $P \in \tilde{\gamma}$, we can sweep Γ with γ so that at any point of time during the sweep, the curves remain non-piercing and P remains within $\tilde{\gamma}$.*

Proof. Let us call a sweeping operation valid if it does not move P out of $\tilde{\gamma}$. Note that vertex bypassing operations are always valid since we can implement them in such a way that P is not moved out of $\tilde{\gamma}$. The only possible sweeping operations that may move P out of $\tilde{\gamma}$ are therefore digon/triangle bypassing.

If we can always find a valid sweeping operation, we can continue applying them eventually ending up in a situation where no pair of curves intersects in the interior of $\tilde{\gamma}$. By Lemma 6, each sweep operation maintains non-piercing. Thus, for contradiction assume that there




■ **Figure 8** If a vertex becomes bypassable after we bypass L , the corresponding bypassability curve can be modified to a bypassability curve for u in the original arrangement.

exists an arrangement of curves so that there are no valid sweep operations. Among such examples let Γ be the *simplest* in the following sense: it lexicographically minimizes (m, n) where m is the number of curves in Γ that lie in the interior of $\tilde{\gamma}$ and n is the number of vertices in $\mathcal{A}(\Gamma)$ that lie in $\tilde{\gamma}$ (including those that lie on γ). If no pair of curves in Γ intersects more than twice in $\tilde{\gamma}$, then by the result of Snoeyink et al., (Theorem 14), there is at least one valid sweeping operation.

The result of Snoeyink et al., can also be applied if some curve α intersects γ in more than 2 points, but no pair of curves intersects more than twice in $\tilde{\gamma}$, since we can treat each segment of α in $\tilde{\gamma}$ as a distinct curve and apply their result. We may thus assume that no curve intersects γ in more than 2 points. Since we assumed that no sweep operations are valid, there must be a pair of curves in $\Gamma \setminus \{\gamma\}$ that intersect 3 or more times in $\tilde{\gamma}$ and therefore form a negative lens inside $\tilde{\gamma}$. Recall that a negative lens is a digon that is not contained in regions bounded by the curves defining it. Let L be a minimal negative lens (i.e., it does not contain another negative lens) in $\tilde{\gamma}$. We can assume that L visible from γ , as otherwise we can bypass it to obtain a simpler non-sweepable arrangement. Suppose that L is defined by the curves α and β . By minimality of Γ , such a lens should exist. Let u and v be the two intersection points of α and β on the boundary of L so that α is oriented from u to v and β is oriented from v to u . Suppose now that we bypass L . The vertices u and v are *lost* as a result of the bypassing of L . We claim that no vertex becomes bypassable as a result of bypassing L . The only vertices that can potentially become bypassable are those that lie in L . Let x_α be a vertex on the boundary of L and that becomes bypassable. We assume without loss of generality that it lies on the modified curve α . Its bypassability curve then *arrives* at x_α from the left of the modified curve α (i.e., from the interior of the new $\tilde{\alpha}$) which means that such a curve must pass arbitrarily close to either u or v . Assume that it passes close to u (see Figure 8a), the other case being analogous. We can then modify it to terminate at u (as shown in Figure 8b) so that it arrives at u from the left of both α and β (i.e., from the interior of both $\tilde{\alpha}$ and $\tilde{\beta}$) in the original arrangement. This shows that u was bypassable before we bypassed the lens L - which by assumption was not the case. Similarly the vertices on the curve β cannot become bypassable as a result of bypassing L . An analogous argument also shows that vertices lying in the interior of L (in the original arrangement) cannot become bypassable after L is bypassed.

We now claim that the new cells created as a result of bypassing L cannot be digon/triangle cells on γ . The only cells that could potentially be digon/triangle cells after bypassing L are those that contain one of the points $u, v, u_\alpha, u_\beta, v_\alpha$ or v_β . The other cells may have α replaced by β (or vice versa) on their boundary as result of bypassing L but this does not change their shape or whether they are on γ . Since they were not digon/triangle cells on γ before, they are not digon/triangle cells after bypassing L .



■ **Figure 9** If a cell with u_α is a digon/triangle cell, then it was bypassable in L .

The cells containing u or u_α cannot be a digon/triangle cell on γ cell since either of those situations would require that the point p shown in Figure 9 (formally p is the first intersection point of α and γ that we arrive at if we start at u on the original curve α and walk along it in the direction opposite to its orientation) is a vertex of the digon/triangle cell. However, that implies that in the original arrangement, we could construct a bypassability curve for u as follows: start at a point on γ arbitrarily close to p that lies to the left of α and follow α closely until u and terminate there - thus arriving at u from the left of both α and β (see Figure 9). This implies that u was bypassable in the original arrangement contradicting our assumptions. By symmetry, the cells containing any of the other points u_β, v, v_α and v_β also cannot be digon/triangle cells. This shows that u must have been bypassable before we bypassed L , contradicting our assumption that this was not the case. ◀

► **Remark.** While our result states that we can sweep by continuously *shrinking* $\tilde{\gamma}$ to a specified point P in $\tilde{\gamma}$, the proof can be easily modified to show that we can sweep by *expanding* $\tilde{\gamma}$ so that a specified point Q remains outside $\tilde{\gamma}$. If Q is chosen to be a point at infinity, then this shows that we can sweep by expanding γ to infinity retaining the property that the curves are non-piercing. The only change required in the proof is that a vertex q defined by curves α and β is said to be bypassable if there is a visibility curve τ from a point on γ to q s.t. the interior of τ lies outside $\tilde{\gamma}, \tilde{\alpha}$, and $\tilde{\beta}$.

7 A simpler proof of the theorem of Snoeyink and Hershberger

In this section, we give a shorter alternative proof of the result of Snoeyink and Hershberger stated in Theorem 14. While our proof follows their general framework, we partition into cases differently and the analysis of the cases is relatively simpler.

Let Γ be an arrangement of 2-intersecting curves with sweep curve γ , and let P be a specified point in $\tilde{\gamma}$. Our goal is to sweep the arrangement of Γ with γ so that the arrangement remains 2-intersecting throughout the process and P remains inside $\tilde{\gamma}$ until no curve in $\Gamma \setminus \gamma$ intersects $\tilde{\gamma}$. We do not use the operation of bypassing a visible vertex, as this violates the 2-intersecting property of the curves. We restrict ourselves to the remaining sweeping operations, namely: (i) taking a loop on a curve that lies entirely in $\tilde{\gamma}$, (ii) bypassing a digon cell, and (iii) bypassing a triangle cell.


A sweeping operation is said to be valid if it does not move P out of $\tilde{\gamma}$. A cell in the arrangement is said to be *sweepable* if it is a digon/triangle cell on γ that does not contain P . Otherwise, it is non-sweepable. If at any point in time, there is a sweepable cell, i.e., there is a valid sweeping operation, we say that the arrangement is *sweepable*. Otherwise, we say that it is non-sweepable.

We will show that the arrangement of any set of 2-intersecting curves with a sweep curve γ and a point $P \in \tilde{\gamma}$ is sweepable. For contradiction, assume that this is not true, and among arrangements that are non-sweepable consider the *simplest* one in the following sense: it

lexicographically minimizes the tuple (c, n, ℓ) where c is the number of curves in Γ lying entirely in the interior of $\tilde{\gamma}$, $n = |\Gamma|$ and ℓ is the number of intersection points (among curves in Γ) lying in $\tilde{\gamma}$ (including those on γ).

Broad idea. In the restricted setting where each pair of curves in $\Gamma \setminus \{\gamma\}$ intersects at most once inside $\tilde{\gamma}$, we can show that the arrangement is sweepable. Therefore, a simplest non-sweepable arrangement must have at least one pair of curves that intersect twice (i.e., they form a digon) in $\tilde{\gamma}$. In this case, we carefully modify the arrangement by bypassing a minimal digon or a minimal triangle (defined in Section 5) to obtain a simpler non-sweepable arrangement, thus arriving at a contradiction.

We start with the following definition that is borrowed from the paper of Snoeyink and Hershberger [21]. Let T be a triangle defined by curves α and β on γ . We say that T is a *half-triangle* with edge β if the side of the triangle defined by β does not intersect any curve in Γ in its interior. Figure 10 shows a half-triangle with edge β .



■ **Figure 10** The figure shows a half-triangle T with edge β containing a triangle cell T' .

► **Lemma 13.** *Let Γ be a set of 2-intersecting curves with sweep curve γ s.t. all curves in $\Gamma \setminus \{\gamma\}$ intersect γ , and any pair of curves in $\Gamma \setminus \{\gamma\}$ intersect at most once in $\tilde{\gamma}$. Then, there is a sweepable cell on γ , i.e., there is either a digon/triangle cell on γ that does not contain P .*

The proof of the main result in this section (Theorem 14) follows the structure of Theorem 1, i.e., proof for non-piercing regions. Unlike in the non-piercing case however, since the curves are 2-intersecting, no pair of curves form a negative lens. Let L be a minimal lens in a simplest non-sweepable arrangement Γ defined by curves $\alpha, \beta \in \Gamma$. We split the proof into two cases: In the first case, there is a curve $\delta \in \Gamma \setminus \{\alpha, \beta\}$ intersecting L . In this case, we show that we can modify the arrangement to obtain a simpler non-sweepable arrangement, thus arriving at a contradiction. If no curve in $\Gamma \setminus \{\alpha, \beta\}$ intersects L , then we split the proof into further sub-cases: Since Γ was assumed to be a simplest non-sweepable arrangement, bypassing L results in a sweepable arrangement. Thus, there is either a digon/triangle cell on γ . If the a newly created cell is a digon cell on γ , we show that removing the curve forming the digon cell results in a simpler non-sweepable arrangement. If we create a triangle cell, we require a slightly more elaborate modification to obtain a simpler non-sweepable arrangement. Thus, in all cases, we obtain a contradiction, and hence there is always either a digon/triangle cell on γ .

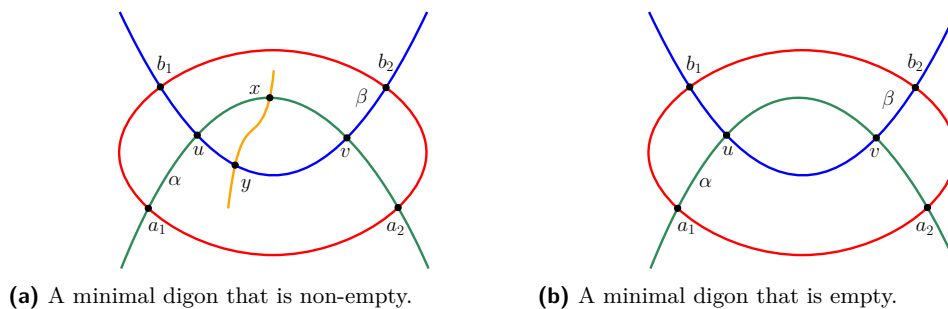
► **Theorem 14** ([21]). *Given any family of 2-intersecting curves Γ , a curve $\gamma \in \Gamma$ and a point $P \in \tilde{\gamma}$, Γ can be swept by γ using the following operations: i) passing a digon cell, ii) passing a triangle cell and iii) taking a loop, so that at any point in time during the sweep, the family of curves remain 2-intersecting, and P lies in the interior of $\tilde{\gamma}$.*

45:12 Sweeping Arrangements of Non-Piercing Regions in the Plane


Proof. Suppose the statement is not true. Consider a simplest non-sweepable family Γ with sweep curve γ i.e., the simplest family for which the statement does not hold. Every curve in this family must intersect $\tilde{\gamma}$ as otherwise we can remove it from Γ to obtain a simpler non-sweepable family. If there is a curve α lying entirely in $\tilde{\gamma}$ it cannot be visible from γ , as otherwise we could take a loop on α (while still keeping P within $\tilde{\gamma}$) to obtain a simpler non-sweepable family - since this decreases the number of regions contained in γ . If α is not visible from γ , $\Gamma' = \Gamma \setminus \{\alpha\}$ is a simpler non-sweepable arrangement contradicting the minimality of Γ . If the curves in $\Gamma \setminus \{\gamma\}$ pairwise intersect at most once in $\tilde{\gamma}$, then by Lemma 13, the arrangement is sweepable.

Therefore, in the arrangement Γ , we can assume that (i) each curve in $\Gamma \setminus \{\gamma\}$ intersects γ twice, and there is a pair of curves intersecting twice in $\tilde{\gamma}$, i.e., they form a digon in $\tilde{\gamma}$, (ii) each minimal lens in $\tilde{\gamma}$ is visible from γ , and (iii) there is either no digon/triangle cell on γ , or there is exactly one digon/triangle C cell in the arrangement, and C contains the point P .

Let L be a minimal digon contained in $\tilde{\gamma}$ and suppose that L is defined by the curves $\alpha, \beta \in \Gamma \setminus \{\gamma\}$ and has vertices u and v . L must be visible from γ , as otherwise by Lemma 8, bypassing L results in a simpler non-sweepable arrangement. Let a_1 and a_2 be the intersection points of α with γ s.t. the points a_1, u, v, a_2 lie in cyclic order along α . Similarly, let b_1 and b_2 be the intersection points of β and γ so that b_1, u, v, b_2 appear in cyclic order along β . See Figure 11. Let $u_\alpha, u_\beta, v_\alpha$ and v_β be the reference points for bypassing lens L . We now split the proof into two cases depending on whether L is *non-empty*, i.e., there is a curve $\delta \in \Gamma \setminus \{\alpha, \beta, \gamma\}$ that intersects L , or is *empty*, i.e., there is no such curve.



■ **Figure 11** The two cases in the proof. Either L is empty or non-empty.



■ **Figure 12** The operations when L is non-empty.

Case 1: L is non-empty. The proof of this case requires the following technical lemma.

▷ **Claim 15.** There exists a curve $\delta \in \Gamma$ that intersects L and such that intersections of any of the other curves with δ within $\tilde{\gamma}$ lie in L . In other words, there are no intersection points on δ that lie in $\tilde{\gamma}$ but outside L .

Since Γ was not sweepable, all cells on γ excluding the cell containing P have at least 4 vertices. Now, consider the arrangement $\Gamma' = \Gamma \setminus \{\delta\}$, where δ is the curve guaranteed by Claim 15. We claim that Γ' remains non-sweepable. Note that the arc of δ between d_2 and y separates two cells C_1 and C_2 . Either one of them contains P , or both cells have at least four vertices. Thus, the new cell created by merging these two cells upon removal of δ is non-sweepable. An analogous statement holds for the arc of δ between d_1 and x . Finally note that the remaining new cells created as a result of removing the portion of δ between x and y are not on cells on γ since they lie within the lens L . Thus, Γ' is a simpler non-sweepable arrangement contradicting the minimality of Γ .

Case 2. L is empty. In this case again, we obtain a contradiction by finding a simpler non-sweepable arrangement. Since we assumed that Γ is a simplest non-sweepable arrangement, bypassing L must create a sweepable cell on γ .

Let D_u be the cell in Γ contained in $\tilde{\gamma}$ with u as a vertex, and has portions of the arcs $[a_1, u]$ and $[b_1, u]$ as two sides. Similarly, let D_v be the cell containing vertex v , and portions of the arcs $[a_2, v]$ and $[b_2, v]$ as sides. By Lemma 9, the only digon/triangle cells on γ in the new arrangement can be the cells that contain one of the reference points $u_\alpha, u_\beta, v_\alpha$, or v_β (See Figure 5). Since L was empty, there is a single cell containing u_α and v_α . Let C_α denote this cell. Similarly, let C_β denote the cell containing the reference points u_β and v_β .

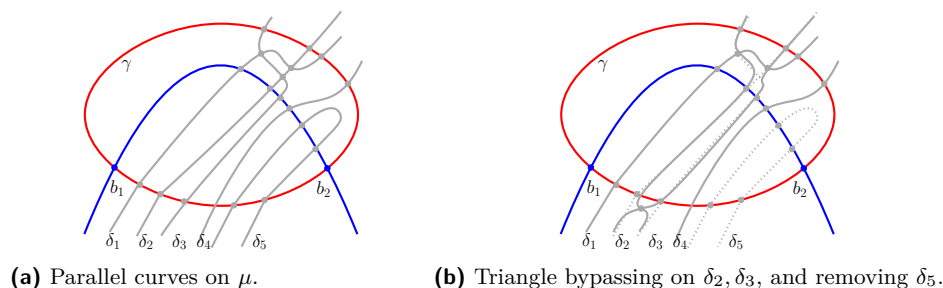
We treat C_α and C_β independently. Let C denote one of these regions. If C is non-sweepable, we don't do anything. If C is a digon cell not containing P , we throw away the corresponding curve (α or β). If C is a triangle cell not containing P , we do a more elaborate modification of the arrangement described below.

We first need some terminology. Let R be a digon bounded by the sweep curve γ and another curve μ with vertices b_1 and b_2 . A *proper chord* λ of R is an arc of another curve δ which intersects the boundary of R exactly once on μ and once on γ . We say that a sequence of proper chords $\lambda_1, \dots, \lambda_k$ are parallel if for $i = 1, \dots, k-1$, the region of R between consecutive chords λ_i and λ_{i+1} is a four-sided cell in the arrangement that does not contain P . For each i , let δ_i be the curve containing λ_i and let Δ_μ denote the sequence $\Delta_\mu = (\delta_1, \dots, \delta_k)$.

▶ **Lemma 16.** *It is possible to modify the curves in $\Delta_\mu = (\delta_1, \dots, \delta_k)$, possibly discarding some of them, so that i) the modified arrangement is as simple, or simpler than Γ , ii) the new digon/triangle cells created in the modified arrangement in $\tilde{\gamma}$ have either b_1 or b_2 as a vertex, and iii) there are no sweepable digon/triangle cells in $\tilde{\gamma}$ on μ other than those having b_1 or b_2 as one of the corners.*


Figures 13 show the modification of the curves in Δ_μ .

Handling triangle cells C_α and/or C_β : Suppose that C_α is a triangle cell T that does not contain P , defined by the curves α, γ and δ_1 . Let R be the digon defined by α and γ that contains T . Without loss of generality let a_1, c_1, d_1 denote the three vertices of T , where c_1 and d_1 are the intersection points of δ_1 with γ and α , respectively. We first apply Lemma 16 with $\mu = \alpha$, R , and we let Δ_α be a maximal sequence of parallel curves in R starting with δ_1 . Let $\Delta'_\alpha = (\delta'_1, \dots, \delta'_l)$ denote the sequence of parallel curves after the



■ **Figure 13** Parallel curves on μ , and operations to remove triangles and digons on μ .

application of Lemma 16. By Lemma 16, the possible sweepable triangle cells in the new arrangement are those in $\tilde{\gamma}$ that contain a_1 or a_2 as a vertex, i.e., the cells containing reference points x_1, y_1, x_2 or y_2 (See Figure 14a). The cell containing reference point y_1 cannot be a sweepable cell it either contains P , or there is at least one more curve other than γ defining this cell that does not intersect α in $\tilde{\gamma}$. Similarly, the cell containing the reference point x_2 cannot be a sweepable cell, as this would imply that the cell in Γ containing x_2 was already sweepable. Therefore, only the cells containing reference points x_1 or y_2 can be sweepable triangle cells. Observe that α, γ and δ'_ℓ form a minimal triangle with base δ'_ℓ . By Corollary 11, bypassing this minimal triangle yields a simpler non-sweepable arrangement of 2-intersecting curves such that the two cells with reference points x_1 and y_1 merge to create a non-sweepable cell, and the cell containing reference point y_2 becomes a 4-sided cell. See Figure 14b. Thus, we obtain a simpler non-sweepable arrangement. If C_β is a triangle cell, we treat it similarly. Thus, in all cases, we obtain a simpler non-sweepable arrangement, contradicting the assumption that Γ was a simplest non-sweepable arrangement. Hence, we can apply one of the sweep operations, and by Corollary 7, each sweep operation preserves the 2-intersecting property of the curves. ◀



■ **Figure 14** Bypassing a minimal triangle.

References

- 1 Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. Coloring hypergraphs defined by stabbed pseudo-disks and abab-free hypergraphs. *SIAM Journal on Discrete Mathematics*, 34(4):2250–2269, 2020.
- 2 Eyal Ackerman, Balázs Keszegh, and Dömötör Pálvölgyi. Coloring delaunay-edges and their generalizations. *Computational Geometry*, 96:101745, 2021.
- 3 Pankaj K Agarwal, Eran Nevo, János Pach, Rom Pinchasi, Micha Sharir, and Shakhar Smorodinsky. Lenses in arrangements of pseudo-circles and their applications. *Journal of the ACM (JACM)*, 51(2):139–186, 2004.

- 4 Bentley and Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Transactions on computers*, 100(9):643–647, 1979.
- 5 Sarit Buzaglo, Rom Holzman, and Rom Pinchasi. On s -intersecting curves and related problems. In *Proceedings of the twenty-fourth annual Symposium on Computational geometry*, pages 79–84, 2008.
- 6 Timothy M. Chan and Elyot Grant. Exact algorithms and apx-hardness results for geometric packing and covering problems. *Comput. Geom.*, 47(2):112–124, 2014. doi:10.1016/J.COMGEO.2012.04.001.
- 7 Timothy M. Chan, Elyot Grant, Jochen Köneemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In Yuval Rabani, editor, *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 1576–1585. SIAM, 2012. doi:10.1137/1.9781611973099.125.
- 8 Bernard Chazelle and Herbert Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *Journal of the ACM (JACM)*, 39(1):1–54, 1992.
- 9 Man-Kwun Chiu, Stefan Felsner, Manfred Scheucher, Felix Schröder, Raphael Steiner, and Birgit Vogtenhuber. Coloring circle arrangements: New 4-chromatic planar graphs. *European Journal of Combinatorics*, page 103839, 2023.
- 10 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>.
- 11 Herbert Edelsbrunner. *Algorithms in combinatorial geometry*, volume 10. Springer Science & Business Media, 1987.
- 12 Herbert Edelsbrunner and Leonidas J. Guibas. Topologically sweeping an arrangement. *Journal of Computer and System Sciences*, 38(1):165–194, 1989. doi:10.1016/0022-0000(89)90038-X.
- 13 Stefan Felsner and Klaus Kriegel. Triangles in euclidean arrangements. *Discrete & Computational Geometry*, 22:429–438, 1999.
- 14 Sarel 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 Chaya Keller, Balázs Keszegh, and Dömötör Pálvölgyi. On the number of hyperedges in the hypergraph of lines and pseudo-discs. *Electron. J. Comb.*, 29(3), 2022. doi:10.37236/10424.
- 16 Rom Pinchasi. A finite family of pseudodiscs must include a “small” pseudodisc. *SIAM J. Discret. Math.*, 28(4):1930–1934, 2014. doi:10.1137/130949750.
- 17 Franco P Preparata and Michael I Shamos. *Computational geometry: an introduction*. Springer Science & Business Media, 2012.
- 18 Rajiv Raman and Saurabh Ray. Constructing planar support for non-piercing regions. *Discret. Comput. Geom.*, 64(3):1098–1122, 2020. doi:10.1007/S00454-020-00216-W.
- 19 Rajiv Raman and Saurabh Ray. On the geometric set multicover problem. *Discret. Comput. Geom.*, 68(2):566–591, 2022. doi:10.1007/S00454-022-00402-Y.
- 20 Manfred Scheucher. *Points, lines, and circles: some contributions to combinatorial geometry*. PhD thesis, Dissertation, Berlin, Technische Universität Berlin, 2019, 2020.
- 21 Jack Snoeyink and John Hershberger. Sweeping arrangements of curves. In Kurt Mehlhorn, editor, *Proceedings of the Fifth Annual Symposium on Computational Geometry, Saarbrücken, Germany, June 5-7, 1989*, pages 354–363. ACM, 1989. doi:10.1145/73833.73872.
- 22 Kasturi R. Varadarajan. Weighted geometric set cover via quasi-uniform sampling. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 641–648. ACM, 2010. doi:10.1145/1806689.1806777.