Sweeping Arrangements of Curves

Jack Snoeyink
Stanford University

John Hershberger
DEC Systems Research Center

Abstract

We consider arrangements of curves that intersect pairwise in at most k points. We show that a curve can sweep any such arrangement and maintain the k-intersection property if and only if k equals 1 or 2. We apply this result to an eclectic set of problems: finding boolean formulæ for polygons with curved edges, counting triangles and digons in arrangements of pseudocircles, and finding extension curves for arrangements. We also discuss implementing the sweep.

1 Introduction

When is it important that lines in an arrangement are straight, or that circles in an arrangement are circular? Of course, it depends on the questions one asks. Many, such as, "What is the minimum number of triangles in an arrangement of lines," do not depend on straightness, but only on the fact that two lines intersect in a single point. [25] One can answer such questions for more general arrangements of pseudolines. Branko Grünbaum, in a lecture entitled "The importance of being straight," points out that because there are arrangements of pseudolines that cannot be stretched to lines, there are questions in which straightness is crucial. He says that we cannot yet answer these because "most of our tools and methods are general (or vague and imprecise?) enough to apply to the case of pseudolines." [20]

In this paper, we look at sweeping arrangements of curves with intersection restrictions. Before we go further, we define *curves, arrangements,* and *sweeping,* and look at reasons to study these objects.

The curves that we consider in this paper lie in the Euclidean plane or on the sphere, are smooth, have no self-intersections, and are endless (either closed or bi-infinite). Any two curves intersect in a finite number of points, at which they cross.

A set of curves Γ has the *k-intersection property* if every two of them intersect in at most k points. If any two curves of Γ intersect in exactly k points, then Γ has the *exact k-intersection property.*

This topological or combinatorial restriction on the intersection of curves is different from the restrictions used in the field of computer graphics. Computer-aided design systems usually place algebraic restrictions on curves; for example, they may require that all curves be lines, conic sections, or cubic plane curves (the components of cubic splines). Natural families of algebraic curves satisfy the k-intersection property for some k, however: lines have the 1-intersection property, vertical parabolas are 2-intersecting curves, general conics are 4-intersecting curves, and cubic plane curves are 9-intersecting. The topological restriction is more general in the sense that if a property holds or an algorithm works for k-intersecting curves, then it will apply to any family of algebraic curves with the k-intersection property.

When calculating arrangements using finite precision arithmetic, one usually cannot preserve straightness. Greene and Yao [18] and Milenkovic [28] have given algorithms that do preserve the 1-intersection property for line arrangements.

Finally, Sharir and other researchers in computational geometry have used intersection restrictions in applying the theory of Davenport-Schinzel sequences to curves [1, 11, 23, 33, 34]. This theory has been an important tool in the analysis of algorithms that deal with algebraic curves.

Unfortunately, our positive results extend only to curves with the 2-intersection property; we cannot say much about cubic plane curves.

A finite set of curves Γ partitions a surface into three types of maximal connected regions: *vertices* are regions contained in two or more curves of Γ, *edges* are contained in only one curve, and *faces* are connected regions contained in no curves of Γ. We call this partition the *arrangement* of Γ. An arrangement is *simple* if no three curves share a common point. We deal primarily with simple arrangements in this paper; we will note where our statements apply to non-simple arrangements. Figure 1 illustrates these definitions.

The names given to the sets of the partition suggest that the arrangement of Γ is a graph embedded in the plane. It makes an unusual graph—some edges are infinite rays, and many edges can connect a pair of vertices. If a curve of Γ does not intersect any other curve, we have an edge with no vertices. With these differences in mind, however, it should not cause confusion to think of the arrangement as a planar graph.

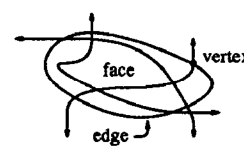

Fig. 1: A simple arrangement

In the past, one studied arrangements of lines and planes in recreational mathematics and because of their relationships to configurations of points and to certain convex polytopes. Grünbaum collected many results and conjectures on arrangements of lines and curves in the plane in his 1972 monograph [22]. Other early results are contained in [19] and [21]. Grünbaum's terminology differs slightly from ours. He discusses "arrangements of pseudolines," "arrangements of curves," and "weak arrangements of curves"; we call them arrangements of curves with the exact 1-, exact 2-, and 2-intersection properties, respectively. We shall call 1-intersecting curves pseudolines and 2-intersecting curves pseudocircles. This means that two pseudolines (or two pseudocircles) are not required to intersect in 1 point (or 2 points).

Recently, researchers in computational geometry have found numerous applications for line arrangements in algorithms for geometry and graphics (see Edelsbrunner's book [7]). They have also considered arrangements of curves with intersection conditions. Edelsbrunner et al. [8] apply Davenport-Schinzel sequences to prove generalizations of the horizon theorem for arrangements of lines [5, 10] and to construct such arrangements incrementally in nearly quadratic time. McKenna and O'Rourke [27] independently proved and used the horizon theorem for the case of pseudocircles.

Sweeping is important both as a paradigm for developing graphical and geometric algorithms, and as tool for use in mathematical proofs. The underlying idea is to determine properties of a collection of objects in a space of dimension d by looking at a series of consecutive $(d-1)$ dimensional slices. Sweeping converts a static problem into a dynamic problem of lower dimension.

As examples of sweep algorithms in the literature, consider the problem of finding the intersections of n lines or segments in the plane. Shamos and Hoey [32] showed how to detect an intersection in $O(n \log n)$ time by sweeping the plane with a line. Bentley and Ottman [2] extended their ideas and developed a practical algorithm to report all K intersections between n segments in $O((n+K)\log n)$ time. Mairson and Stolfi [26] used a sweeping line to report all K intersections between two sets of segments, each of which has no self-intersections, in $O(n \log n + K)$ time. They also applied their techniques to segments of x-monotone curves.

If a straight line sweeps the plane to find the intersection points of a set of segments, then it encounters the points in sorted order in the direction perpendicular to the sweep. In such a sweep, the logarithmic factors in the time complexity seem inescapable. For infinite lines Edelsbrunner and Guibas [9] showed that we can avoid the logarithmic factor by sweeping the plane with a pseudoline. Their algorithm runs in optimal $O(n^2)$ time. Chazelle and Edelsbrunner [4] recently developed an algorithm to report all K segment intersections in $O(n \log n + K)$ time that also sweeps the plane with a pseudoline.

In the next section we will define precisely what it means to sweep arrangements of curves that have the k-intersection property, but let's first think about what condition we would want on the sweeping curve c. In designing algorithms, we would like the curves of our arrangement to have few intersections with c, since we must keep track of each intersection. In proofs dealing with a family of k-intersecting curves, we would like c to fit into the family.

In section 3 we prove the sweeping theorem, which says that we can always sweep an arrangement of pseudolines or pseudocircles starting from any curve in the arrangement, and cannot always sweep k-intersecting curves for $k \geq 3$.

Theorem 3.1 (Sweeping theorem) *Let Γ be a finite set of bi-infinite curves (and closed curves, if $k > 1$) in the plane or sphere that have the k-intersection property. Let c be a curve of Γ. If $k \in \{1, 2\}$ then we can sweep Γ starting with c and maintain the k-intersection property. If $k > 2$ then arrangements exist that cannot be swept.*

In the next section we define local operations by which a sweeping curve can advance. For pseudolines ($k = 1$) and pseudocircles ($k = 2$) we can always apply one of the local operations. For $k > 2$ we indicate how to construct an arrangement that cannot be swept by a given curve.

In section 4 we present two applications of the sweep theorem. We use the pseudoline case to extend the work of Dobkin et al. [6]; we find a short boolean formula to describe a polygon with curved edges. We use the pseudocircle case to find a relationship between the minimum number of digons and triangles in an arrangement of 2-intersecting curves. This is related to Grünbaum's conjecture 3.7 [22].

Section 5 defines another type of sweep in the plane: sweeping a double wedge. We use this sweep to prove an extension theorem for arrangements that includes Levi's lemma [25] as a special case.

Theorem 5.1 (Extension theorem) *Let Γ be a finite set of curves with the k-intersection property, and P be a set of $k+1$ points, not all on the same curve. If $k \in \{1, 2\}$ then there is a curve that contains the points of P and has the k-intersection property with respect to Γ. If $k > 2$ then arrangements and point sets exist such that any curve through the points violates the k-intersection property.*

Section 6 points out that we can easily implement the local operations used in the proof of the sweeping theorem when we know the arrangement. When we do not, we can apply the ideas of Edelsbrunner and Guibas [9] to sweep a set of pseudolines using linear space. We leave sweeping pseudocircles as an open problem.

The operations that we define in section 2 can also sweep arrangements of ($k > 2$)-intersecting curves if we do not require the sweep to have the k-intersection property. But is there some intersection property that is maintained by such a sweep? For algorithmic purposes, it would be satisfactory to have a function of k (and not of n) bound the number of points of intersection the sweep has with any curve.

2 Definitions for sweeping

In the beginning of this section we define sweeping as a continuous process; later we will see that we can carry it out in discrete steps. We also define notation that we use in the proofs of section 3.

Let c be an endless curve in the sphere or Euclidean plane with an orientation, which we depict as left to right. Since c is smooth and has no self-intersections, it divides the plane or sphere into two connected components; one with boundary oriented clockwise, the other counterclockwise. We say that the component with the counterclockwise orientation is *above* c and the other component is *below*. A curve γ *lies above* c if the component above c contains γ. When we sweep, we handle these two components separately.

To sweep the component of the plane or sphere above c, we want to move c continuously to infinity if the component is unbounded, or to a point if it is bounded. Formally, we say that a family of curves indexed by the positive reals, $C = \{c_\alpha\}_{\alpha \geq 0}$, *covers the component above* c if

- The curve $c_0 = c$,
- Every point above c (except one, if the component above c is bounded) lies on exactly one curve c_α, and
- The curve c_β is above c_α whenever $\beta > \alpha$.

To say that we can *sweep the component above* c means that a family of curves C exists that covers the component above c.

Suppose c is a member of a set of k-intersecting curves, Γ. We can *sweep the arrangement of* Γ *above* c if there is a family of curves C that covers the component above c, and the set $\Gamma \cup C$ has the k-intersection property. Finally, we say that we can *sweep* Γ *starting from* c if we can sweep above c and below c.

Though we employ this continuous definition to prove negative results in section 3.3, it is unwieldy for proving positive results. In the literature, sweeping is investigated as a discrete process by identifying events where the intersection between the sweep and the plane changes significantly. One can usually interpolate between two consecutive curves of a discrete sweep in a natural way to obtain a continuous family that covers the plane or sphere.

In our case, the oriented sweeping curve c intersects the other curves of Γ in some order; this order changes when c passes a vertex of the arrangement or changes the number of its intersections with some curve. If we choose these changes as our events and find the curves of the sweep where the events occur, then we obtain a discrete set of curves. We can interpolate between two consecutive curves with the help of the Schönflies theorem [31]: Map adjacent curves c and c' to parallel lines. Map the curve segments between c and c' to non-intersecting line segments in the strip between the lines. The Schönflies theorem implies that we can carry out this mapping by a homeomorphism of the plane or sphere to itself. Under the inverse mapping, the parallel lines in the strip become a family of curves that interpolate between c and c' and intersect the curves of the arrangement in a consistent order. By interpolating between each pair of adjacent curves, we can extend the curves from a discrete sweep to a continuous family that covers the plane.

We will define local operations that allow us to advance the sweeping curve past an intersection point of two curves and to add or remove curves from the set of curves intersected by the sweep, without violating the k-intersection property. But first, let us give names to some of the things that appear along the sweeping curve.

We can represent the portion of the arrangement Γ that lies on and above the sweep c as a graph $G = \langle V, E \rangle$. (We assume that $c \in \Gamma$.) Let A denote the component of the plane or sphere on and above c. The vertices V of G are the points in A where two curves intersect. The edges E are the edges of the arrangement that are contained in A. To avoid having edges that are not adjacent to vertices, we introduce an artificial vertex on any curve that has no intersection points. We denote the edges in E by $\langle u, v \rangle$, where $u, v \in V \cup \{\infty\}$. An edge $\langle u, \infty \rangle$ is an infinite ray.

Assume that we are sweeping above c. A point p on a curve γ is *visible from an edge* e of c if there is a face above c whose boundary contains both p and the edge e. In particular, p *is visible from* c if p

lies on a face immediately above c.

Suppose that m of the vertices of G appear on the sweeping curve c. We number these vertices with the integers $1, \ldots, m$ in order of their appearance along c. In this and the next section, we use the names i, j, k and l for vertices on c, and the names u, v and w for vertices above the sweep.

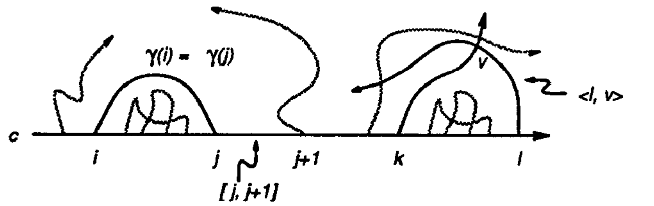

Fig. 2: A hump $\langle i,j \rangle$ and a triangle $\Delta(k,v,l)$ in G

Figure 2 illustrates some other notation that we will use. If the intersection of the curve $\alpha \in \Gamma$ with c forms the vertex i, then we define $\gamma(i) = \alpha$. Some edges of the graph G lie on the curve c—we give these edges special status and denote them with square brackets, for example, $[j, j+1]$, instead of angle brackets. Thus, if we write $\langle i,j \rangle$ we mean the edge from vertex i to vertex j along the curve $\gamma(i)=\gamma(j)$, and not an edge along the curve c.

We also define terms to describe some special configurations in G. A *hump* is an edge $\langle i,j \rangle$ of G with both i and j on the sweeping curve c. The hump $\langle i,j \rangle$ is *empty* if it contains no other curves. Clearly, for an empty hump, $|j-i| = 1$. The hump pictured is not empty. A *triangle* is a pair of edges, $\langle k,v \rangle$ and $\langle l,v \rangle$, with two distinct vertices k and l both on the sweep and one vertex v above the sweep—we denote this triangle by $\Delta(k,v,l)$. The triangle $\Delta(k,v,l)$ is *empty* if it contains no other curves. For an empty triangle, $|l-k| = 1$.

With this notation, we can now define the local operations that move c forward while preserving the k-intersection property. The operations, shown in Figure 3, are:

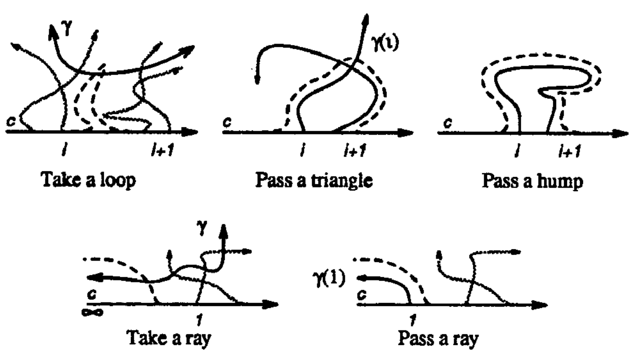

Fig. 3: Operations by which the sweep progresses

1. *Taking a loop.* Let γ be a curve that intersects the sweep in at most $k-2$ points. If an edge e of γ is visible from an edge $[i, i+1]$ of c, then there is a path from a point on e to a point on $[i, i+1]$ that intersects no other curves. The sweep c can advance along this path and intersect the edge e in two places without adding any other intersections—preserving the k-intersection property.

2. *Passing an empty triangle.* If $\Delta(i, v, i+1)$ is an empty triangle, then c can move past the vertex v, interchanging the order of $\gamma(i)$ and $\gamma(i+1)$ along the sweep.

3. *Passing an empty hump.* If the edge $\langle i, i+1 \rangle$ is an empty hump, then c can advance past it and reduce the number of intersections with $\gamma(i)$ by two.

4. *Taking the first ray.* If a bi-infinite curve γ intersects the sweep in fewer than $k-1$ points and is visible as an infinite ray from edge $[\infty, 1]$ of c, then c can move forward and introduce one intersection point with this ray; γ becomes $\gamma(1)$. We can define *taking the last ray* in a similar fashion.

5. *Passing the first ray.* If the leftmost curve $\gamma(1)$ is an infinite ray with no intersections above the sweep c, then c can move past the ray and lose one intersection point with $\gamma(1)$.

If we can apply an operation, we say the sweep can make progress. A discrete sweep terminates when either there are no curves above the sweep or the only curves above the sweep are non-intersecting rays. In these situations, one can continuously sweep above c and maintain the k-intersection property without difficulty.

In a non-simple arrangement, we may encounter a number of triangles with a common vertex v above the sweep. We can handle this either by passing all the triangles simultaneously or by perturbing the curves (either actually or conceptually) to form a simple arrangement. Since our theorems will not restrict how the operations are applied, either of these methods will work.

3 The sweeping theorem

Using the definition of sweeping from the previous section, we now prove:

Theorem 3.1 (Sweeping theorem) *Let Γ be a finite set of bi-infinite curves (and closed curves, if $k > 1$) in the plane or sphere that have the k-intersection property. Let c be a curve of Γ. If $k \in \{1,2\}$ then we can sweep Γ starting with c and maintain the k-intersection property. If $k > 2$ then arrangements exist that cannot be swept.*

We establish the pseudoline and pseudocircle cases, $k = 1$ and $k = 2$, by showing that certain local operations can always make progress. In section 3.3 we present some unsweepable arrangements with $k > 2$.

3.1 Sweeping Pseudolines

To prove the pseudoline case of the sweep theorem, we show that the sweep can always make progress using three of the operations defined in section 2.

Lemma 3.1 *Any arrangement of bi-infinite curves Γ with the 1-intersection property can be swept starting with any curve $\gamma \in \Gamma$ using three operations: passing a triangle, passing the first ray, and taking the first ray.*

Proof: We will see that the sweep can either pass the first ray, take the first ray, or pass some triangle.

Consider the edge $\langle 1, u \rangle$ of the curve $\gamma(1)$ that intersects the sweep first. If $u = \infty$, then we can pass $\gamma(1)$. Otherwise the vertex u comes from the intersection of a curve α_1 with $\gamma(1)$.

Suppose that the curve α_1 does not intersect the sweep and that the curves $\{c, \gamma(1), \alpha_1, \ldots, \alpha_m\}$ border an unbounded face in counterclockwise order, as illustrated in Figure 4. We can show by induction that α_m does not intersect the sweep.

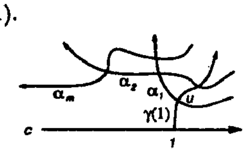

Fig. 4: α_1 does not intersect the sweep

Assume that these curves are oriented in the counterclockwise direction. Then the curve α_{i-1} divides α_i into two pieces, one to the left and one to the right. If α_{i-1} does not intersect the sweeping curve c, then it separates the right piece of α_i from c. The left piece of α_i must remain in the region bounded by $\{c, \gamma(1), \alpha_1, \ldots, \alpha_{i-1}\}$; since vertex 1 is the first vertex on c, the left piece also cannot intersect c. Therefore, if α_1 does not intersect c, then neither does α_m, and c can take the first ray.

Otherwise α_1 intersects the sweep at a point ℓ.

Let us call a triple (i, v, j) a *half-triangle* if $\langle i, v \rangle$ is an edge and $\gamma(i)$ and $\gamma(j)$ intersect at v. No curve crosses the edge $\langle i, v \rangle$ of (i, v, j), so any curve visible in the interval (i, j) crosses c between i and j and $\gamma(j)$ between v and j. By induction on the size of the interval (i, j), we can show that every half-triangle contains an empty triangle.

In the base case, $|i-j|=1$ and $\Delta(i,v,j)$ is an empty triangle. Otherwise a curve β intersects $\gamma(j)$ at w such that $\langle j,w\rangle$ is an edge. Since β is visible from j, it intersects c at $k \in (i,j)$. But now (j,w,k) is a smaller half-triangle, which contains an empty triangle by the induction hypothesis.

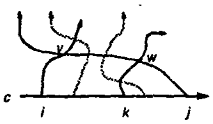

Fig. 5: An empty triangle in a half-triangle

Since α_1 intersects the sweep at ℓ, the triple $(1,u,\ell)$ is a half triangle. The sweep can make progress by passing the empty triangle contained in $(1,u,\ell)$. This establishes the lemma. ∎

3.2 Sweeping Pseudocircles

To establish the pseudocircle ($k=2$) case of the sweep theorem we use operations that change an even number of intersections. Once again we show that the sweep can always advance.

Lemma 3.2 *Any arrangement of curves Γ with the 2-intersection property can be swept starting from any curve $\gamma \in \Gamma$ by using three operations: passing a triangle, passing a hump, and taking a loop.*

We first prove that a bi-infinite curve can sweep an arrangement of 2-intersecting curves curves in the the Euclidean plane. We assume that a counterexample exists and then derive a contradiction. At the end of this section, we show that a closed sweeping curve can sweep a sphere or the Euclidean plane.

For each arrangement in the plane with a bi-infinite sweeping curve we can form a graph G of the portion on or above the sweep. We define the *size* of an arrangement with a sweeping curve to be the pair $(|V|,n)$; that is, the number of vertices on and above the sweep followed by the number of curves. Now, choose a set to 2-intersecting curves Γ, including an oriented bi-infinite sweeping curve c, such that c cannot make progress and the size of the arrangement is lexicographically minimum. We prove a sequence of lemmas about the structure of this arrangement and wind up with a contradiction. This contradiction implies that a bi-infinite sweeping curve can always sweep an arrangement of 2-intersecting curves.

Lemma 3.3 *Removing any curve γ from the arrangement Γ allows the sweep to progress.*

Proof: By minimality of the arrangement. ∎

Lemma 3.4 *All curves in Γ intersect the sweeping curve c.*

Proof: Since the sweep cannot make progress by taking a new curve, any curve γ that does not intersect c is not visible from c. But removing such a curve γ does not change c's ability to make progress. Therefore, all curves in the smallest counterexample intersect the sweeping curve c. ∎

Lemma 3.5 *We can choose a minimum counterexample such that all of the curves of Γ are rays to infinity below the sweep c.*

Proof: Map the bi-infinite sweeping curve c to a line c' by a continuous mapping of the plane onto itself, cut all curves below c' and extend them perpendicularly to infinity, then apply the inverse mapping. When this procedure is applied to a minimum counterexample, neither the arrangement above the sweep nor the size is affected. Thus, it remains a minimum counterexample. ∎

From now on, we assume that our unsweepable arrangement, Γ, is chosen in accordance with lemma 3.5.

Lemma 3.6 *The arrangement Γ cannot contain a hump or a triangle.*

Proof: The arrangement certainly cannot contain an empty triangle or hump; otherwise the sweep could make progress. Suppose, however, that Γ contains vertices i and j on the sweep, with $i<j$, that are the ends of a hump or a triangle. Since none of the curves that intersect the sweep between $i+1$ and $j-1$ intersect the hump or triangle, the smaller arrangement, $\Gamma' = \{\gamma(i+1),\gamma(i+2),\ldots,\gamma(j-1)\}$, lies completely inside the hump or triangle. But, by minimality, the sweep can make progress in the arrangement Γ'. This contradicts the fact that Γ is unsweepable and establishes the lemma. ∎

Let's pause a moment and consider what these first four lemmas tell us. We have an arrangement Γ in which the sweep cannot progress. If we remove any curve γ, then some operation applies. Lemmas 3.4 and 3.6 tell us that, after removing γ, there must be an empty hump or triangle where there were none before. We also have a notational convenience from lemma 3.4; every curve in Γ can be denoted $\gamma(i)$ for some vertex i on the sweep.

Next we consider configurations in which removing a curve leaves a hump or triangle. By finding contradictions to lemma 3.6 or to the minimality of Γ, we can show that most configurations cannot occur.

Lemma 3.7 *For any vertex j on the sweep, the removal of the curve $\gamma(j)$ cannot leave an empty hump or empty triangle with vertices i and k satisfying $i<j<k$.*

Proof: By case analysis. We sketch the idea: Lemma 3.6 implies that the curve $\gamma(j)$ intersects the triangle or hump that is formed if $\gamma(j)$ is removed. That intersection point forms a new empty triangle. To break up the new triangle, $\gamma(j)$ must reenter it—but in doing so, the curve $\gamma(j)$ creates another triangle that cannot be avoided. ∎

Lemma 3.8 *The removal of a curve $\gamma(k)$ cannot leave a hump.*

Proof: Suppose that the arrangement of $\Gamma - \{\gamma(k)\}$ contains the hump $\langle i,j\rangle$. Consider the arrangement Γ' with only this hump removed: $\Gamma' = \Gamma - \{\gamma(i)\}$. By lemma 3.3, the sweep can make progress in Γ'. But $\gamma(i)$ intersects only $\gamma(k)$; therefore, in Γ', the curve $\gamma(k)$ contributes an edge to an empty hump or triangle—let $\langle k,v\rangle$ be that edge. In the original arrangement Γ, however, $\gamma(i)$ intersects $\gamma(k)$ at some vertex u between k and v. Thus $\Delta(i,u,k)$ is a triangle in Γ, contradicting lemma 3.6. ∎

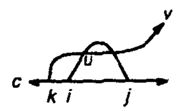

Fig. 6: $\gamma(k)$ cannot leave a hump

If the removal of a curve γ does not allow us to move past a hump (lemma 3.8) or take an edge (lemma 3.4) then we must have an empty triangle, $\Delta(i,v,j)$. Since γ cannot intersect the sweep between i and j (lemma 3.7), we can denote the triangle in the arrangement $\Gamma - \{\gamma\}$ by $\Delta(i,v,i+1)$. The next lemma will restrict the way γ can break up this triangle.

Lemma 3.9 *Suppose that the arrangement $\Gamma - \{\gamma\}$ has an empty triangle $\Delta(i,v,i+1)$. Then, in the original arrangement Γ, the face above the edge $[i,i+1]$ has v as a vertex.*

Proof: Suppose that v is not a vertex of the face above the edge $[i,i+1]$. Then the curve γ must "hide" v from that portion of the sweep by intersecting both legs, $\langle i,v\rangle$ and $\langle i+1,v\rangle$ of the triangle $\Delta(i,v,i+1)$. Lemma 3.7 implies that γ intersects both legs once or both legs twice.

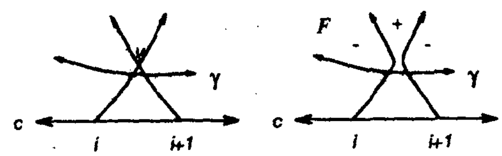

Fig. 7: The curve γ intersects the legs once. We untangle the triangle.

In both cases we can derive a contradiction by "untangling" the triangle as shown in figure 7. Untangling eliminates the vertex v, forming a smaller arrangement that can be swept. We omit the proof to save space. ∎

We have shown that removing any curve γ from the minimum counterexample Γ must leave an empty triangle and that γ must intersect only one leg of that triangle. If removing γ leaves $\Delta(i, v, i+1)$ then we say that the edge $[i, i+1]$ is *responsible* for γ. Notice that each edge of the sweep is responsible for at most one curve.

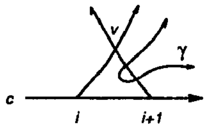

Fig. 8: $[i, i+1]$ is responsible for γ

In order to have an edge of the sweep responsible for every one of the n curves in the set Γ, we need at least $n+1$ vertices on the sweep. Thus, there is a curve that intersects the sweep twice. We call a curve a *loop with respect to* c, or, more simply, a *loop*, if it intersects c twice and is connected in the above component of the plane.

The presence of a loop is critical to our argument. If we were sweeping an infinite cylinder, we would not always have a loop; in fact, there are unsweepable arrangements of curves that have the two-intersection property and have no loops. Figure 9 shows an arrangement in which the sweeping circle, c, cannot make progress. The existence of a loop is a property of the plane that prevents the sweep from becoming stuck.

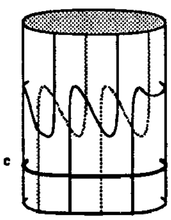

Fig. 9: Cannot be swept

We omit the proof of the final lemma for the bi-infinite case.

Lemma 3.10 *There exists a curve γ whose removal does not leave a triangle.*

Proof: We define a nesting for loops and look at an innermost loop γ. By looking at γ and the interval resonsible for γ we can find a smaller arrangement in which the sweep can make progress if and only if it can make progress in the original. But this contradicts the minimality of the original arrangement. ∎

Lemma 3.3 states that if we remove one curve from the smallest counterexample, then some operation applies. Lemmas 3.4 and 3.8, however, say that the operations of taking a curve or passing a hump will never apply after removing one curve. Lemma 3.10 proves the existence of a curve whose removal does not leave a triangle. This contradiction proves that any arrangement of bi-infinite, 2-intersecting curves curves in the Euclidean plane can be swept by a curve that maintains the 2-intersection property.

Having established the bi-infinite case, we can use it to prove that a closed curve can always make progress. To sweep a sphere with a closed curve c, simply find a point that lies only on c and remove this point from the plane. Sweeping this punctured sphere is equivalent to sweeping the Euclidean plane with a bi-infinite curve.

To sweep an arrangement in the Euclidean plane with a closed curve, we show how to embed it in an arrangement of 2-intersecting curves on the sphere. The following lemma shows that we can map the plane into a region of a sphere bounded by a 2-intersecting curve ξ. Then we close off the bi-infinite curves outside of ξ to form an arrangement on the sphere with the 2-intersection property.

Lemma 3.11 *The bounded faces of an arrangement of endless curves with the 2-intersection property in the plane can be embedded into a region bounded by a curve in an arrangement of closed 2-intersecting curves in the sphere.*

Proof: Draw an auxiliary curve ξ around the arrangement Γ so that all the vertices and closed curves of Γ are on the interior of ξ and modify the infinite rays so ξ intersects each ray once. Take a solid half-ball H and map ξ to the rim of H so that the interior of the curve ξ is mapped to the rounded surface of the half-ball. Every bi-infinite curve of Γ has two points of intersection with ξ; connect these two points by a chord along the flat surface of the half-ball. These chords turn the bi-infinite curves of Γ into closed curves; we must show that any two closed curves intersect in at most two points.

Suppose two chords, $\overline{ab}$ and $\overline{cd}$, intersect. Let α denote the closed curve including chord $\overline{ab}$, and let β denote the closed curve including chord $\overline{cd}$. The non-chordal portion of β goes from the interior of the closed curve α to the exterior, so it intersects α an odd number of times. But, by the 2-intersection property, it must do so only once. Therefore, closing off the curves of Γ with these chords does not violate the 2-intersection property. ∎

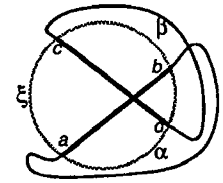

Fig. 10: Closing off curves with chords

Given an arrangement in the Euclidean plane and a closed sweeping curve c, embed the arrangement in the sphere according to lemma 3.11. To sweep the interior of c, simply sweep on the sphere. To sweep the exterior of c, sweep on the sphere until c reaches the auxiliary curve ξ. Then return to the plane, cutting c where it meets ξ and extending it to infinity. Finish the sweep in the plane with this bi-infinite sweeping curve. Thus, a closed sweeping curve can always progress and lemma 3.2 is established.

3.3 Arrangements that cannot be swept

There are arrangements with the 3-intersection property that cannot be swept.

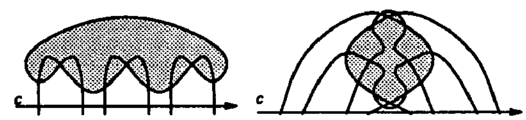

Fig. 11: Unsweepable arrangements of 3-intersecting curves

Consider the arrangements in figure 11. In both of them, the sweeping curve c intersects every curve twice. When the sweep first passes over one of the vertices on the boundary of the shaded region, it intersects one or two curves, each in two additional places—in both instances violating the 3-intersection property. Thus these arrangements are unsweepable.

It is not difficult to construct arrangements with the $(k > 3)$-intersection property that have copies of the arrangements of figure 11 stopping progress in triangles or humps. This establishes the final case of the sweeping theorem.

4 Applications of sweeping

In this section we apply the sweeping theorem to solve two problems: finding boolean formulæ for polygons with curved edges and counting triangles and digons in arrangements of exactly 2-intersecting curves.

4.1 Boolean formulæ for polygons

The lines supporting the edges of a simple polygon define half-spaces in the plane. We can describe the set of points in the interior of the polygon by a boolean formula on these half-spaces. For example, a convex polygon is the AND of the half-spaces defined by its edges. Peterson [30] showed that every polygon with n edges can be represented by a monotone formula that uses each half-space once.

Any bi-infinite curve divides the plane into two half-spaces; we can attempt to find formulæ for polygons with curved edges if these edges are pieces of bi-infinite curves. Define a *curved segment* to be a simply connected portion of a curve. We call the points that bound a segment *vertices* in this section; segments with only one endpoint we call *rays*. A *polygonal chain* is a sequence of curved segments $\{s_1, s_2, \ldots, s_n\}$ such that s_i and s_{i+1} share a common vertex, either s_1 and s_n are rays (and the chain is *open*) or they share a common endpoint (and the chain is *closed*), and no other intersections between segments occur. A vertex v of a segment s divides a bi-infinite curve into two pieces; we call the piece not containing the segment the *extension of s through v*.

Dobkin et al. [6] gave a simple proof that every polygon with n edges has a *Peterson-style formula*—a monotone boolean formula that uses each edge's half-space once. Their proof depends on the fact that every polygonal chain has a *splitting vertex*—a vertex v such that the (straight-line) extensions of the incident edges through v do not intersect the chain. They split the chain at v, recursively find formulæ for the two subchains, and combine them with an AND if v is convex or an OR if v is concave.

We can use the sweep theorem to show that polygonal chains whose segments are portions of pseudolines also have a splitting vertex.

Theorem 4.1 *In any polygonal chain whose segments are connected portions of distinct pseudolines, there is a vertex v such that the extensions of the incident segments through v do not intersect the polygonal chain.*

Proof: Let $S = \{s_1, s_2, \ldots, s_n\}$ be a polygonal chain satisfying the hypothesis.

If S is closed, sweep its arrangement of pseudolines starting from a bi-infinite curve that intersects no curves. Stop when the sweeping curve c first crosses the chain S. Since all segments of S are bounded, the sweep first intersects S by passing a triangle and not by taking or passing the first ray. Hence, the intersection point crossed is a vertex v of the chain, as shown in figure 12a. The incident segments both intersect c, so their extensions through v lie wholly on the swept side of c. The chain S, however, lies on the unswept side, except for v. Therefore the extensions cannot intersect S.

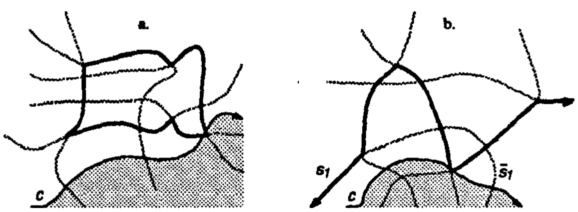

Fig. 12: Sweeping closed and open polygonal chains

If S is open, let $\bar{s}_1$ be the ray that is the extension of s_1 through its vertex. Assume that the rays s_1, $\bar{s}_1$, and s_n appear counterclockwise in this order on the line at infinity; if they are reversed, consider the mirror arrangement. Sweep the arrangement beginning with a bi-infinite curve c that intersects only $\bar{s}_1$ and stop when c intersects the chain S, as shown in figure 12b.

Recall that we sweep by passing triangles and taking and passing the first (and not the last) ray. Since the sweep c intersects $\bar{s}_1$, it cannot cross the ray s_1; ray s_1 blocks c from crossing s_n. Furthermore, c cannot pass $\bar{s}_1$ without intersecting the chain S since $s_1 \in S$. Therefore c first intersects S by passing a triangle, and the extensions of the segments incident to the triangle's apex are separated from the chain S as in the previous case. This proves the theorem. ∎

This theorem, together with the results of [6], proves that polygons bounded by portions of pseudolines have Peterson-style formulæ.

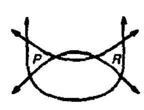

Fig. 13: P has no formula

Polygons bounded by portions of pseudocircles may require auxiliary curves and variables. In figure 13, the polygon R is contained in the same half-spaces as the polygon P—to represent one region without the other, we must separate them with another curve.

4.2 Triangles and digons in arrangements of exactly 2-intersecting curves

We can model lines and pseudolines in the projective plane by great circles on a sphere with oposite points identified. In "Arrangements and Spreads," Grünbaum notes that if we cease to identify opposite points of the sphere, we obtain an arrangement of exactly 2-intersecting curves with no digons (two-sided faces) [22]. He investigates minimum and maximum number of triangles in such arrangements and asks what is the relationship between the number of triangles and number of digons.

We can prove:

Lemma 4.1 *In an arrangement of closed curves with the exact 2-intersection property, let p_i denote the number of i-sided faces. Then* $2p_2 + 3p_3 \geq 4n$.

Proof: Assume the arrangement is on the sphere. We will find four digons or triangles on each curve.

Choose a curve c. Map c and one of the hemispheres it defines to a disk in the plane by the Schönflies theorem. [31] Extend the curves to infinity outside the disk.

Sweep the disk starting with c. Since c already intersects all of the curves, it advances by passing a triangle or hump (a digon). Cut c inside the hump or triangle, and extend the ends to infinity without crossing any other curves. We can still sweep the disk with this bi-infinite curve, so there is another digon or triangle. Similarly, we find two digons or triangles in the other hemisphere defined by c.

Since we find each digon at most two times and each triangle at most three times, we have $2p_2 + 3p_3 \geq 4n$. ∎

Grünbaum conjectures [22, Conj. 3.7] that every digon-free arrangement of exactly 2-intersecting curves has at least $2n-4$ triangles. Specializing our result proves there are at least $4n/3$ triangles—we suspect that his conjecture is closer to the truth.

5 Sweeping wedges

In the previous section our sweeping curves moved forward to sweep the plane. In this section we will see how to sweep by rotation about a point. We use this sweep to prove the extension theorem, which extends arrangements of 1- and 2-intersecting curves by adding new curves that pass through specified points.

Theorem 5.1 (Extension theorem) *Let Γ be a finite set of curves with the k-intersection property, and P be a set of $k+1$ points, not all on the same curve. If $k \in \{1,2\}$ then there is a curve that contains the points of P and has the k-intersection property with respect to Γ. If $k > 2$ then arrangements and point sets exist such that any curve through the points violates the k-intersection property.*

Levi proved the exact $k = 1$ case of the extension theorem in 1926 [25]. Levi's extension lemma, as his result is known, has become an important tool for generalizing properties of arrangements of lines to arrangements of pseudolines. For example, Goodman and Pollack use it, in concert with their "circular sequences" [7, 17], in a series of papers that prove duals of Radon's and Helly's theorems for pseudolines [16], prove that all arrangements of 8 pseudolines are stretchable [13], and define a duality for configurations of points and pseudolines [12, 15] to establish a conjecture of Burr, Grünbaum, and Sloane [3]. They also show that no extension lemma can hold for planes in space [14].

As we mentioned in section 4.2, Grünbaum [22] is interested in arrangements of exactly 2-intersecting curves. He says, "Another open and seemingly hard problem is to find the right analogue for (digon-free) arrangements of curves of Levi's extension lemma. It is well possible that an appropriate result in this direction would lead to solutions of some of the other problems mentioned." The pseudocircle case of the extension theorem finds a 2-intersecting curve, which is not necessarily an exactly 2-intersecting curve—we do not guarantee that our new curve intersects all other curves.

We establish both the pseudoline and pseudocircle cases of the extension theorem by showing that a bi-infinite curve can sweep a double wedge in the Euclidean plane. Let us define such a sweep formally before we give the proofs.

Let c and c' be oriented bi-infinite curves that intersect only at a point p. Curves c and c' divide the plane into four regions bounded by curve segments, as illustrated in figure 14. We call the region where both segments are oriented towards p the *left wedge*, and call the region where both are oriented away from p the *right wedge*. The left and right wedges together comprise the *double wedge*, denoted $wedge(c, c')$. The wedge $wedge(c, c')$ contains the curves c and c'.

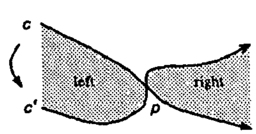

Fig. 14: A double wedge

The segment of c bounding the left wedge is the *left half* of c, that bounding the right wedge is the *right half*. We assume that the left wedge is below c and the right wedge is above, according to the definitions of below and above given in section 2. We want to rotate c continuously, moving downward on the left and upward on the right until c reaches c'.

A family of curves, $\mathcal{C} = \{c_\alpha\}_{0 \le \alpha \le 1}$, indexed by reals from the interval $[0, 1]$, *covers the double wedge* $wedge(c, c')$ if

- The curves $c_0 = c$ and $c_1 = c'$,
- Every point in $wedge(c, c')$ lies on exactly one curve c_α except the point p, which lies on every curve in $wedge(c, c')$, and
- For all $0 \le \alpha < \beta < \alpha' \le 1$, the curve c_β lies in $wedge(c_\alpha, c_{\alpha'})$

Suppose that c and c' form a wedge $wedge(c, c')$ and belong to a set of curves Γ that has the the k-intersection property. We say that c can *sweep the double wedge* $wedge(c, c')$ if there is a family of curves $\mathcal{C}$ that covers $wedge(c, c')$, and the set $\Gamma \cup \mathcal{C}$ has the k-intersection property. As before, we assume that the sweeping curve c is part of the arrangement.

A point q on a curve γ is visible from an interval I of the sweep c, if q and some point $r \in I$ both lie on the boundary of a common face contained in the double wedge $wedge(c, c')$.

We can view the sweeping process in another way. If we remove the point p from the plane, the resulting surface is topologically equivalent to a cylinder. The sweeping curve c and its destination c' each become two bi-infinite curves. Both pieces of c are moving in the same direction (counterclockwise in figure 15). These two pieces together are not allowed to intersect any curve in more than k points; the sweep ends when the pieces of c reach the pieces of c'.

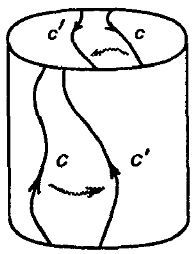

Fig. 15: c on a cylinder

Even though we have defined sweeping as a continuous operation, we again want to perform the sweep in discrete steps. We use the local operations from section 2—modifying the "taking" operations to use the new definition of visibility. We prove the extension theorem in the next three subsections.

5.1 Extension lemma for pseudolines

We do not need to sweep a double wedge to prove the extension theorem for exactly 1-intersecting curves. Levi originally proved his lemma by arguing that a curve that connected two points and had the minimum number of intersections with other curves had the exact 1-intersection property. See [25] or, for a proof in English, [22, Thm. 3.4]. Our proof, which sweeps a double wedge, handles the 1-intersecting curves case and foreshadows the proof the pseudocircle case.

Suppose we have an arrangement of pseudolines and points p and q not on the same pseudoline. If no curve passes through p then we add one to the arrangement by sweeping with a bi-infinite curve until we encounter p. Now, if only one curve passes through p, duplicate it, reverse its orientation and perturb it so that the original and the copy intersect only at p and the point q lies in the double wedge of these curves. Otherwise, q lies in some region bounded by two curves through p. Orient these two curves so that q lies in their double wedge. Lemma 5.1 proves that we can sweep the double wedge containing q and maintain the 1-intersection property. This establishes the $k = 1$ case of the extension theorem.

Lemma 5.1 *Given a finite set of pseudolines Γ that includes two curves, c and c', that define a double wedge, the curve c can sweep $wedge(c, c')$ using the operations of passing a triangle and taking or passing the first or last ray.*

Proof: First apply operations that advance the right half of the sweep c. If c can no longer make progress in the right wedge, then it will be able to advance in the left wedge.

Assume c cannot advance on the right—no operation applies in the right wedge. We perform surgery: Erase everything in the plane except c and the curve segments in the right wedge, and extend these segments to bi-infinite curves without adding any intersections. The only effect of this surgery that is visible from the right half of c is that c' disappears. Lemma 3.1 says that c can make progress in the reduced arrangement by passing triangles and passing and taking the last ray. Since no operation applied in the original arrangement and the curves visible from c did not change, c cannot pass a triangle

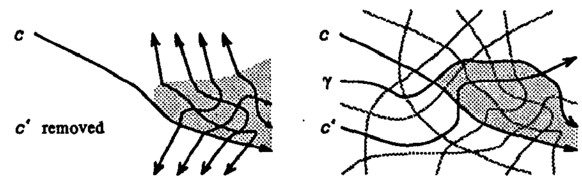

Fig. 16: Surgery on the right wedge finds curve γ

or ray after the surgery. Therefore c can now take a curve γ that intersected the left half of c before the surgery.

Look at the right wedge of $wedge(c, \gamma)$. Since γ is visible from c as a ray in this wedge, every pseudoline that intersects the right half of c also intersects γ above c.

Returning to $wedge(c, c')$, perform the same surgery as above on the left wedge. Lemma 3.1 says that c can advance downward in the reduced arrangement by passing triangles and passing and taking the first ray. As before, if we can pass a triangle or a ray, we can also pass it in the original arrangement. If we can take the ray of a curve α, then α intersects γ below the sweep. Therefore, α cannot intersect the right half of c—if it did, it would intersect γ above the sweep as well. This means that c can take α in the original arrangement.

Therefore, c can advance on the left to sweep the double wedge. ∎

This completes the proof of the pseudoline case.

5.2 Extension theorem for pseudocircles

We will also prove the extension theorem for pseudocircles by sweeping a double wedge. First, however, we show that we need consider only closed curves on the sphere. Given an arrangement of pseudocircles in the Euclidean plane, lemma 3.11 says we can map it to a region $\mathcal{R}$ of a sphere bounded by a pseudocircle and close off all the curves without violating the 2-intersection property. Let p, q, and r denote the three points given by the hypothesis; they too are mapped into $\mathcal{R}$. Assume for the moment that we can find a curve c through p, q, and r on the sphere. The curve c intersects the boundary of $\mathcal{R}$ in at most two points, so the portion inside $\mathcal{R}$, that is $c' = c \cap \mathcal{R}$, is a simply connected curve. This curve c' maps to a 2-intersecting curve in the Euclidean plane.

Now, consider the sphere with an arrangement of closed pseudocircles and three points p, q, and r—for simplicity, assume no curves pass through these three points. If we remove r, then we need to find a bi-infinite curve that passes through p and q. As in the pseudoline case, the sweeping theorem says that we can sweep with a bi-infinite curve until we encounter p. Duplicate the sweep, reverse its orientation, and perturb it so that the original and the copy intersect only at p and the point q lies in the double wedge of these curves. Lemma 5.2 proves that we can sweep the double wedge containing q and establishes the $k = 2$ case of the extension theorem.

Lemma 5.2 *Let Γ be a finite set of pseudocircles, all closed except two, c and c', which define a double wedge. The curve c can sweep $wedge(c, c')$ using the operations of passing a triangle, passing a hump, and taking a loop.*

In some arrangements, lemma 5.3 pinpoints a place where an operation applies. In other arrangements, we find curves that we call crescents. Crescents have regions, which we call their shadows, in which the sweep can advance unless some other crescent prevents it. By looking at the structure of crescents and their shadows, we will see that there is some crescent whose shadow contains no other crescents. The sweeping lemma for pseudocircles, lemma 3.2, is an important tool in this proof.

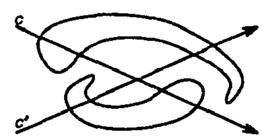

Fig. 17: Two crescents

We call a curve a *crescent* if it is visible from one half of the sweep and intersects the other half in two points. Keep in mind that all curves but c and c' are closed. We will refer to crescents that intersect the left (right) half as above (below) crescents.

Crescents are important because:

Lemma 5.3 *If an arrangement does not have both above and below crescents, then the sweep can make progress.*

Proof: Suppose the arrangement has no above crescents; the other case is similar. Then perform surgery on the right wedge: erase everything from the plane except c and the curve segments in the right wedge, then extend these segments to bi-infinite curves without adding intersections. (See figure 16.)

Lemma 3.2 says that c can now advance by passing a hump or triangle or by taking a loop. If a passing operation applies in the reduced arrangement, then it also applies in the original. On the other hand, no curve visible from the right half of the sweep intersects the left half, so any loop that can be taken in the reduced arrangement can also be taken in the original. Therefore, c can make progress in sweeping the double wedge. ∎

Assume for now that γ is a below crescent—it intersects the right half of the sweep c and is visible from a point p in the left half of c. The hump of γ is the portion of γ above the right half of c. If c could pass the hump of γ on the right, it could subsequently take the loop of γ on the left. We make the following definitions in an attempt to capture the set of curves (or curve segments) that prevent c from passing this hump.

We define the *inner and outer shadows* of γ recursively in the following paragraph. The inner shadow $\mathit{inSh}(\gamma)$ is a set of curve segments, each having an endpoint that is before the hump of γ on the right half of c. The outer shadow $\mathit{outSh}(\gamma)$ is a set of curve segments, each having an endpoint that is after the hump of γ on the right half of c. The shadow $\mathit{Sh}(\gamma)$ is the union $\mathit{inSh}(\gamma) \cup \mathit{outSh}(\gamma)$. The *shadow interval* of Sh, $\mathit{intvl}(\mathit{Sh})$, is the portion of c from the leftmost to the rightmost point of $c \cap \mathit{Sh}$. The ***shadow region*** of Sh is the largest simply connected region bounded by the curve segments of the shadow Sh and the interval $\mathit{intvl}(\mathit{Sh})$.

Let $\mathit{inSh}_0 = \mathit{outSh}_0$ be the hump of γ. We form the set inSh_{i+1} by adding certain curve segments to inSh_i: For every curve α that intersects the right half of c before the interval $\mathit{intvl}(\mathit{inSh}_i)$ and has a point p_α visible from $\mathit{intvl}(\mathit{inSh}_i)$, we add the curve segment from p_α to the intersection point with the right half of c nearest $\mathit{intvl}(\mathit{inSh}_i)$. If no curves are added, then we terminate the recursion and set $\mathit{inSh}(\gamma) = \mathit{inSh}_i$. Since a finite number of curves intersect c, the recursion eventually terminates.

We use a similar recursion for the outer shadow $\mathit{outSh}(\gamma)$, replacing "before" by "after" in the previous paragraph.

The *leading curve* of inner shadow inSh_i (outer shadow outSh_i) is the curve whose intersection with c is the first (last) point of $\mathit{intvl}(\mathit{inSh}_i)$ ($\mathit{intvl}(\mathit{outSh}_i)$). The set of leading curves for an inner shadow inSh_i (outer shadow outSh_i) is the set $\{\pi_0, \pi_1, \ldots, \pi_i\}$, where π_j is the leading curve for inSh_j (outSh_j).

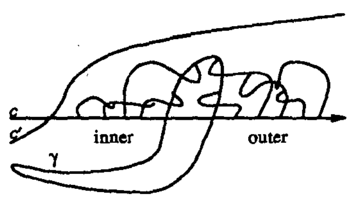

Fig. 18: The shadows

We have made these definitions for below crescents. For above crescents, we make similar definitions involving the left half of the sweep. The only difference is that the outer shadow precedes the inner shadow so that the inner shadow is closer to the vertex of the double wedge $\mathit{wedge}(c, c')$.

The way curves nest inside each other is important in our proofs. We say that a curve α ***nests*** with a non-crescent curve β of the inner shadow $\mathit{inSh}(\gamma)$ (outer shadow $\mathit{outSh}(\gamma)$) if both points of the intersection $\alpha \cap c$ lie inside (outside) the closed curve β. We say that a curve α *nests inwardly* with the below crescent γ if both points of $\alpha \cap c$ lie in the interval of c between the hump of γ and p, the point from which γ is visible below c. The curve α *nests outwardly* with γ if both points of $\alpha \cap c$ lie outside of the hump of γ and the interval from the hump to p.

Curves that contribute curved segments to a shadow have nesting properties that are described in the following lemma whose proof.

Lemma 5.4 *Let β be a curve whose segment β' is added in forming the shadow $\mathit{inSh}_i(\gamma)$ (or $\mathit{outSh}_i(\gamma)$). Let α be a curve that intersects the sweep c, but not in the shadow interval. If α intersects β', then α nests with β and with $\pi_0, \pi_1, \ldots, \pi_{i-1}$, the set of leading curves of $\mathit{inSh}_{i-1}(\gamma)$ (or $\mathit{outSh}_{i-1}(\gamma)$).*

Corollaries to this lemma prove the definition of shadows is unambiguous and reveal some structure of the shadows.

Corollary 5.5 *The construction of the inner and outer shadows of γ is well defined.*

Proof: When adding the segment of a curve α to an inner shadow inSh_i, we add the segment "from p_α to the intersection point with the right half of c nearest $\mathit{intvl}(\mathit{inSh}_i)$." Lemma 5.4 says that α nests with γ. Therefore, both points of $\alpha \cap c$ either precede or follow $\mathit{intvl}(\mathit{inSh}_i)$ and the nearest one is well defined. ∎

Corollary 5.6 *Let α be a curve with a point p_α that is visible from the shadow interval $\mathit{intvl}(\mathit{Sh}(\gamma))$. Either α is a crescent, does not intersect c, or connects p_α to c within the shadow region of $\mathit{Sh}(\gamma)$.*

Proof: Assume α intersects c and is not a crescent. If α is first visible from but does not intersect $\mathit{intvl}(\mathit{inSh}_i)$ (or $\mathit{intvl}(\mathit{outSh}_i)$), then lemma 5.4 says that α nests with the leading curve of the shadow. Thus α will be included in inSh_{i+1} (or outSh_{i+1}). ∎

Now we can show how one crescent can prevent the sweep from advancing in the shadow of another. We say that an above or below crescent A *stops* a below or above crescent B, $A \oslash B$, if A is visible from the shadow interval $\mathit{intvl}(\mathit{Sh}(B))$.

Lemma 5.7 *If no crescent stops a crescent γ, then the sweep can advance in the shadow region of $\mathit{Sh}(\gamma)$.*

Proof: Suppose that no operation applies in $\mathit{intvl}(\mathit{Sh}(\gamma))$. As usual, we perform surgery: erase every curve outside the shadow region except c and extend all curve segments to infinity without adding intersections.

The sweep lemma, lemma 3.2, says that c can make progress in the reduced arrangement. If c could advance by passing a hump or triangle, then it could do so in the original arrangement. Therefore it can take a loop of a curve α. Since α is not a crescent, corollary 5.6 implies that α does not intersect c. Thus c can take α in the original arrangement. ∎

To prove that there is a crescent not stopped by any other crescent, we show that the $\oslash$ relation is acyclic. We omit the proof due to space constraints.

Lemma 5.8 *The $\oslash$ relation is acyclic.*

Proof: If cycles exist, we can show that the crescents of a minimal length cycle must nest deeper and deeper. The most deeply nested crescent, however, cannot stop any other crescent of the cycle. This contradiction shows that no cycle exists. ∎

This lemma, combined with lemma 5.7, shows that we can make progress even when crescents exist. Thus, the proof of lemma 5.2 is complete.

5.3 Arrangements with no extension curves

Figure 19 shows that when we specify $k+2$ points in pseudoline ($k = 1$) or pseudocircle ($k = 2$) arrangements, an extension curve may not exist.

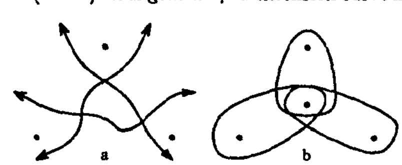

Fig. 19: No extension curve

In figure 19a, we want to draw a new curve through the three points, so one point lies between the other two. But then one of the pseudolines separates the middle point from the adjacent points—the new curve intersects this pseudoline at least twice.

In figure 19b, we can assume that the new curve through these four points is closed. The center point then has two adjacent points and one opposite point on this closed curve. One of the pseudocircles separates the center and its opposite from their adjacent points. The new curve intersects this pseudocircle in at least four points. Since the pseudocircles of figure 19b also satisfy the 3-intersection property, the above argument proves that not all arrangements of 3-intersecting curves with four points have an extension curve. We will show how to convert this to an arrangement of bi-infinite curves in lemma 5.9.

Lemma 5.9 *For every $k \geq 3$, there is an arrangement of k-intersecting curves and $k+1$ points such that any curve through these points violates the k-intersection property.*

Proof: If $k > 3$ is even, consider arrangements of four curves constructed after the scheme in figure 20. The figures on the left show arrangements with the 4 and 6-intersection property. Those on the right indicate how to form arrangements for other even k—curve segments are drawn to represent the open curves that surround their perimeter.

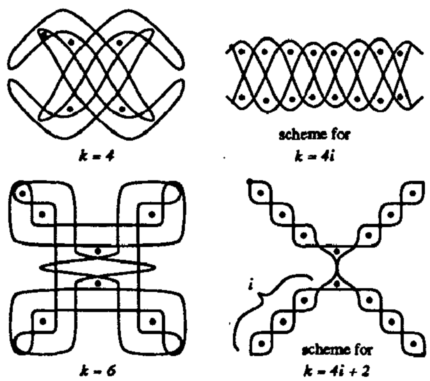

Fig. 20: Arrangements for even k

We can connect one pair of points by a (curved) segment that crosses curves 2 times. Any other segment that connects two points crosses the curves at least 4 times. To connect all $k+1$ points requires $4k+2$ crossings, or an average of $k+1/2$ crossings per curve. Thus, any extension curve through the points intersects one of the curves in more than k points.

Since the extension must cross each curve an even number of times, it intersects one of the curves in $k+2$ points. Thus, the arrangements of figure 20 also show $(k+1)$-intersecting curves with no extension curves.

For those who are not happy with closed curves satisfying an odd intersection property, we can replace each closed curve by three bi-infinite curves as shown in figure 21. An extension curve can extend infinitely in only two of the three, so we can close one without crossing the extension curve. This reduces the problem to the previous case. ∎

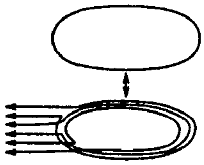

Fig. 21: Bi-infinite curves

6 Algorithms for sweeping curves

If we know the graph that represents the arrangement, then we can easily implement the local operations defined in section 2. The only complication could be dealing with separate connected components of the graph. If we sweep pseudolines or pseudocircles, lemmas 5.1 and 5.2 say that we can apply any operation to make progress—we may even be able to apply operations in parallel.

Since an arrangement of n curves can have quadratic size, it is important to ask whether one can perform the sweep when the arrangement is not known explicitly. Edelsbrunner and Guibas showed that this could be done efficiently for line arrangements [9]. Let's look at their ideas and try to use them to sweep arrangements of pseudolines and pseudocircles.

6.1 Implementing a pseudoline sweep

When sweeping an arrangement of lines, Edelsbrunner and Guibas can start with a curve that intersects every line—they don't need to use taking or dropping operations. Thus, their problem is to recognize empty triangles. For this purpose they keep two data structures that they call upper and lower horizon trees.

The horizon trees could also be called *envelope trees* because for any interval $[i, i+1]$ on the sweep, one tree encodes the lower envelope of the lines that intersect the sweep before the interval, specifically the lines $\gamma(1), \ldots, \gamma(i)$, and the other encodes the upper envelope of the lines $\gamma(i+1), \ldots, \gamma(n)$. (Unfortunately, the upper horizon tree encodes the lower envelopes.) In this section, we want to concentrate on the envelope properties, so we will refer to the trees as lower and upper envelope trees.

Any empty triangle that occurs in both envelope trees is truly empty. The lower envelope tree certifies that no line cuts it from above and the upper that no line cuts it from below. Since we can pass triangles in any order, we need only one tree: one can prove that in the upper envelope tree (lower horizon tree) the uppermost or first triangle cannot be cut by a line above it. Thus if we always pass the upper triangle in the upper envelope tree, we can sweep the plane. Overmars and Welzl noticed this fact in the dual [29].

For lines, and also for pseudolines, the envelope trees have linear size—once a curve crosses into the envelope, it cannot leave. The initial trees are easy to construct. Simply add the curves in increasing order along the sweep to build the lower envelope tree, and in decreasing order to build the upper envelope tree. The time to update the trees can be related to the horizon complexity of the lines of the arrangement (thus the original names) so it amortizes to constant time per triangle.

We can now modify the method of Edelsbrunner and Guibas to use a single envelope tree to sweep pseudolines in linear space and in time proportional to the size of the arrangement. Initially we need to know the order along the sweep c of the curves that intersect c. We also assume that we know the ordering around the line at infinity. With this information, we can take rays until we reach a curve that already intersects the sweep—say it is γ intersecting the sweep at i. That is, $\gamma = \gamma(i)$.

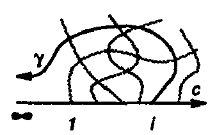

Fig. 22: Sweeping curves $\gamma(1), \ldots, \gamma(i)$

We find the upper envelope tree (lower horizon tree) for the curves $\Gamma' = \{\gamma(1), \ldots, \gamma(i)\}$. The first triangle in this tree is an empty triangle in the arrangement Γ'. Every curve that is visible in the interval $[\infty, i]$ of the sweep c must intersect c at or before i to avoid intersecting γ twice; therefore the first triangle is an empty triangle in Γ and c can pass it. The sweep c can continue to advance until c drops the curve γ. If the sweep is not yet complete, then c returns to taking curves.

6.2 Open problems for $(k > 1)$-intersecting curves

Applying these ideas to sweep pseudocircles seems more difficult. There are two obvious complications: Since pseudocircles can be closed, we may have to sweep several connected components of the graph of their arrangement. The fact that each curve can intersect the sweep twice is certain to complicate the description if not the algorithm.

A special case that avoids these two complications is the case in which the sweep intersects every pseudocircle once. Since passing a hump and taking a loop change two intersections, lemma 3.2 says that an arrangement of such curves can be swept by passing triangles. An algorithm for this case might have application to computing skewed projections [24, 27].

There are complications even in this case. The envelope trees can have size $\Omega(n^2)$, so we don't want to store them explicitly. We need some structure that can store them compactly and allow efficient updates when two curves change order. We leave this as an open problem.

Another open problem involves sweeping arrangements of k-intersecting curves, for $k > 2$. We can sweep any arrangement by first applying our local operations until we get stuck, then taking curves, violating the k-intersection property, until other operations apply. In an extreme case, if the sweep takes all visible edges, then clearly it can make progress. What intersection property does the sweeping curve satisfy if we sweep with this procedure? How can such a procedure be implemented if the arrangement is not known explicitly? These questions are important for practical applications to sweeping curves.

References

[1] P. Agarwal, M. Sharir, and P. Shor. Sharp upper and lower bounds on the length of Davenport-Schinzel sequences. Technical Report 332, Department of Computer Science, New York University, 1987.

[2] J. L. Bentley and T. A. Ottman. Algorithms for reporting and counting intersections. *IEEE Transactions on Computers*, C–28:643–647, 1979.

[3] S. A. Burr, B. Grünbaum, and N. J. A. Sloane. The orchard problem. *Geometriæ Dedicatia*, 2:397–424, 1974.

[4] B. Chazelle and H. Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. In *Proceedings of the 29th IEEE Symposium on Foundations of Computer Science*, pages 590–600, 1988.

[5] B. Chazelle, L. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25:76–90, 1985.

[6] D. Dobkin, L. Guibas, J. Hershberger, and J. Snoeyink. An efficient algorithm for finding the CSG representation of a simple polygon. *Computer Graphics*, 22(4):31–40, 1988. Proceedings of SIGGRAPH '88.

[7] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer Verlag, Berlin, 1987.

[8] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir. Arrangements of curves in the plane—topology, combinatorics, and algorithms. In *Proceedings of the 15th International Colloquium on Automata, Languages and Programming*, 1988.

[9] H. Edelsbrunner and L. J. Guibas. Topologically sweeping an arrangement. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing*, pages 389–403, 1986.

[10] H. Edelsbrunner, J. O'Rourke, and R. Seidel. Constructing arrangements of lines and hyperplanes with applications. *SIAM Journal on Computing*, 15:341–363, 1986.

[11] H. Edelsbrunner, J. Pach, J. T. Schwartz, and M. Sharir. On the lower envelope of bivariate functions and its applications. In *Proceedings of the 28th IEEE Symposium on Foundations of Computer Science*, pages 27–37, 1987.

[12] J. E. Goodman. Proof of a conjecture of Burr, Grünbaum, and Sloane. *Discrete Mathematics*, 32:27–35, 1980.

[13] J. E. Goodman and R. Pollack. Proof of Grünbaum's conjecture of the stretchability of certain arrangements of pseudolines. *Journal of Combinatorial Theory, Series A*, 29:385–390, 1980.

[14] J. E. Goodman and R. Pollack. Three points do not determine a (pseudo-) plane. *Journal of Combinatorial Theory, Series A*, 31:215–218, 1981.

[15] J. E. Goodman and R. Pollack. A theorem of ordered duality. *Geometriæ Dedicatia*, 12:63–74, 1982.

[16] J. E. Goodman and R. Pollack. Helly-type theorems for pseudoline arrangements in P^2. *Journal of Combinatorial Theory, Series A*, 32:1–19, 1984.

[17] J. E. Goodman and R. Pollack. Semispaces of configurations, cell complexes of arrangements. *Journal of Combinatorial Theory, Series A*, 37, 1984.

[18] D. H. Greene and F. F. Yao. Finite-resolution computational geometry. In *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science*, pages 143–152, 1986.

[19] B. Grünbaum. *Convex Polytopes*. John Wiley & Sons, London, 1967.

[20] B. Grünbaum. The importance of being straight. In *Proceedings of the Twelfth Biennial Seminar of the Canadian Mathematical Congress on Time Series and Stochastic Processes; Convexity and Combinatorics*, pages 243–254, 1970.

[21] B. Grünbaum. Arrangements of hyperplanes. In *Proceedings of the Second Louisiana Conference on Combinatorics, Graph Theory and Computing*, pages 41–106, 1971.

[22] B. Grünbaum. *Arrangements and Spreads*, volume 10 of *Regional Conference Series in Mathematics*. American Mathematical Society, Providence, RI, 1972.

[23] S. Hart and M. Sharir. Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes. *Combinatorica*, 6:151–177, 1986.

[24] J. W. Jaromczyk and M. Kowaluk. Skewed projections with an application to line stabbing in R^3. In *Proceedings of the 4th Annual ACM Symposium on Computational Geometry*, pages 362–370, 1988.

[25] F. Levi. Die Teilung der projektiven Ebene durch Gerade oder Pseudogerade. *Ber. Math. Phys. Kl. Sächs. Akad. Wiss. Leipzig.*, 78:256–267, 1926.

[26] H. G. Mairson and J. Stolfi. Reporting line segment intersections. In R. Earnshaw, editor, *Theoretical Foundations of Computer Graphics and CAD*. NATO ASI Series, F40, Springer Verlag, 1988.

[27] M. McKenna and J. O'Rourke. Arrangements of lines in 3-space: A data structure with applications. In *Proceedings of the 4th Annual ACM Symposium on Computational Geometry*, pages 371–380, 1988.

[28] V. J. Milenkovic. *Verifiable Implementations of Geometric Algorithms Using Finite Precision Arithmetic*. PhD thesis, Carnegie-Mellon University, Pittsburg, Penn., 1988. Technical Report CMU-CS-88-168.

[29] M. H. Overmars and E. Welzl. New methods for computing visibility graphs. In *Proceedings of the 4th Annual ACM Symposium on Computational Geometry*, pages 164–171, 1988.

[30] D. Peterson. Halfspace representation of extrusions, solids of revolution, and pyramids. SANDIA Report SAND84-0572, Sandia National Laboratories, 1984.

[31] D. Rolfsen. *Knots and Links*. Publish or Perish, Berkeley, 1976.

[32] M. I. Shamos and D. Hoey. Geometric intersection problems. In *Proceedings of the 17th IEEE Symposium on Foundations of Computer Science*, pages 208–215, Houston, 1976.

[33] M. Sharir, R. Cole, K. Kedem, D. Leven, R. Pollack, and S. Sifrony. Geometric applications of Davenport-Schinzel sequences. In *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science*, pages 77–86, 1986.

[34] A. Wiernik and M. Sharir. Planar realization of non-linear Davenport-Schinzel sequences by segments. *Discrete & Computational Geometry*, 3:15–47, 1988.