

## On Arrangements of Jordan Arcs with Three Intersections per Pair\*

Herbert Edelsbrunner,<sup>1</sup> Leonidas Guibas,<sup>2,3</sup> John Hershberger,<sup>3</sup>  
Janos Pach,<sup>4,5</sup> Richard Pollack,<sup>4</sup> Raimund Seidel,<sup>6,7</sup>  
Micha Sharir,<sup>4,8</sup> and Jack Snoeyink<sup>2</sup>

<sup>1</sup> University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

<sup>2</sup> Stanford University, Stanford, CA 94305, USA

<sup>3</sup> DEC Systems Research Center, 130 Lytton Avenue, Palo Alto, CA 94301, USA

<sup>4</sup> New York University, 251 Mercer Street, New York, NY 10012, USA

<sup>5</sup> Hungarian Academy of Sciences, Pf. 127, H-1364 Budapest, Hungary

<sup>6</sup> IBM Almaden Laboratory, 650 Harry Road, San Jose, CA 95120-6099, USA

<sup>7</sup> University of California at Berkeley, Berkeley, CA 94720, USA

<sup>8</sup> Tel Aviv University, 69978 Tel Aviv, Israel

**Abstract.** Motivated by a number of motion-planning questions, we investigate in this paper some general topological and combinatorial properties of the boundary of the union of  $n$  regions bounded by Jordan curves in the plane. We show that, under some fairly weak conditions, a simply connected surface can be constructed that exactly covers this union and whose boundary has combinatorial complexity that is nearly linear, even though the covered region can have quadratic complexity. In the case where our regions are delimited by Jordan arcs in the upper halfplane starting and ending on the  $x$ -axis such that any pair of arcs intersect in at most three points, we prove that the total number of subarcs that appear on the boundary of the union is only  $\Theta(n\alpha(n))$ , where  $\alpha(n)$  is the extremely slowly growing functional inverse of Ackermann's function.

\* The first author is pleased to acknowledge the support of Amoco Fnd. Fac. Dev. Comput. Sci. 1-6-44862 and National Science Foundation Grant CCR-8714565. Work on this paper by the fourth and seventh authors has been supported by Office of Naval Research Grant N00014-87-K-0129, by National Science Foundation Grant NSF-DCR-83-20085, and by grants from the Digital Equipment Corporation and the IBM Corporation. The seventh author in addition wishes to acknowledge support by a research grant from the NCRD—the Israeli National Council for Research and Development. The fifth author would like to acknowledge support in part by NSF grant DMS-8501947. Finally, the eighth author was supported in part by a National Science Foundation Graduate Fellowship.

## 1. Preliminaries

Recent work in planning a collision-free motion for a robot has given rise to a number of interesting questions in combinatorial geometry. For example, Kedem *et al.* [KL\*] posed and answered the following question: Suppose we are given  $n$  closed Jordan curves<sup>1</sup> in the plane. Denote the curves by  $\gamma_1, \gamma_2, \dots, \gamma_n$ , and their interiors by  $K_1, K_2, \dots, K_n$ , respectively. If we assume that each pair of curves  $\gamma_i, \gamma_j$ , with  $i \neq j$ , intersect each other in at most two points, how many intersection points total can there be on the boundary of  $K = \bigcup_{i=1}^n K_i$ ? The answer to this question was important in the analysis of an efficient algorithm investigated in [KL\*] for planning a collision-free translational motion of a convex polygon amidst polygonal obstacles.

If we allow as many as four intersection points per pair of curves, the number of intersection points on the boundary of the union of the interiors of our curves can easily become  $\Theta(n^2)$ , as Fig. 1 shows. Notice that this behavior is attained even when our closed Jordan curves are triangles and that, asymptotically speaking, a quadratic number of intersections is the worst possible—as long as there is a fixed bound on the number of intersections of any particular pair of curves. In [KL\*] it is shown that, under the assumption of at most two intersections per pair of curves, the total number of intersection points on the boundary of  $K$  can be at most  $O(n)$ . This linear bound allowed the authors to prove the efficiency of a simple algorithm for the stated motion-planning problem.

How many intersection points can there be on the boundary of  $K$  if we allow at most three intersection points per pair of curves? Under the assumption that all intersections have to be proper (i.e., transverse) this question is improperly stated. Any pair of closed Jordan curves must have an even number of proper intersections, so any pair of curves intersecting in at most three points intersects in fact in at most two points. In order to make better sense of the question, let us modify the set-up as follows.

We let each  $\gamma_i$  be a Jordan arc, i.e., an open curve without self-intersections, starting and ending on the  $x$ -axis, and otherwise lying fully above the  $x$ -axis. Two such arcs  $\gamma_i$  and  $\gamma_j$ , with  $i \neq j$ , can have any number of proper intersections; in particular, one or three intersections are possible. We find it convenient to

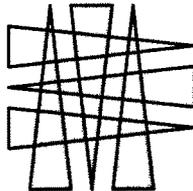


Fig 1. A quadratic number of intersections.

<sup>1</sup> A Jordan curve is one without self-intersections.

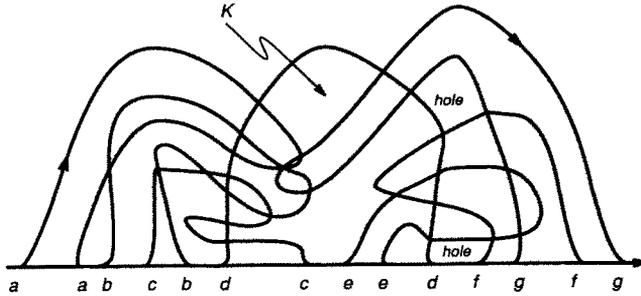


Fig. 2. Three-intersecting Jordan arcs in the upper halfplane.

extend each arc  $\gamma_i$  to be a bi-infinite curve by adding to it the portions of the  $x$ -axis from  $-\infty$  to its left endpoint, and then from its right endpoint to  $+\infty$ . We also orient each such bi-infinite arc from left to right, in the obvious way. Every arc thus extended subdivides the plane into two open regions. The one lying to the right of  $\gamma_i$  we call the interior of  $\gamma_i$  and denote by  $K_{\gamma_i}$ , abbreviated  $K_i$ ; note that each  $K_i$  contains the lower halfplane. We also find it convenient to assume that the  $x$ -axis itself is always one of our arcs.

So now our question becomes, given  $n$  such Jordan arcs in the upper halfplane, with at most three intersections per pair above the  $x$ -axis (the three-intersection condition), how many intersection points total can there be on the boundary of  $K = \bigcup_{i=1}^m K_i$ ? We call this the *three-intersection* problem. See Fig. 2 for an example.

In this paper we prove that the answer to this question is  $\Theta(n\alpha(n))$ , where  $\alpha(n)$  is an extremely slowly growing function of  $n$  that is the functional inverse of Ackermann's function. The same function arises in the theory of Davenport-Schinzel sequences developed by Sharir [HS], [S], and in its many geometric applications. We also show that the same bound holds for the combinatorial complexity of the exterior boundary of the union of any  $n$  "halfspaces" whose bounding curves are bi-infinite and have the property that each pair of these curves intersects in at most three points.

The three-intersection problem is a generalization of several problems whose solution is already known. For example, assume each arc  $\gamma_i$  is the graph of a partially defined positive function  $y = \gamma_i(x)$  (whose domain of definition is the interval between the two endpoints of  $\gamma_i$ ), and that any two distinct such graphs intersect in at most one point (the pseudoline condition). Complete each  $\gamma_i$  to be a totally defined function by drawing two rays steeply descending from its endpoints until they encounter the  $x$ -axis. Then it is easily checked that any pair of these total functions intersect in at most three points. But in this case the boundary of  $K$  (in the previous notation) is the *upper envelope* of the functions  $\gamma_i(x)$ ; by Davenport-Schinzel theory [A], [HS], the number of intersections of the  $\gamma_i$  on the boundary of  $K$  is at most  $\lambda_3(n) = O(n\alpha(n))$ . A more special case of this result can be obtained as follows. Take a collection of  $n$  segments  $e_1, e_2, \dots, e_n$  so that each  $e_i$  is nonvertical and lies above the  $x$ -axis. For each  $i$ , let  $\gamma_i$  be the arc obtained by connecting the two endpoints of  $e_i$  to the  $x$ -axis by

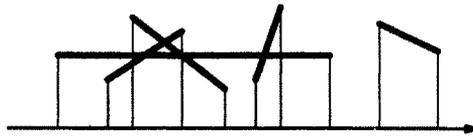


Fig. 3. The upper envelope of segments.

two vertical segments (see Fig. 3). Clearly, each pair of these arcs  $\gamma_i$  intersects in at most three points. The upper envelope of the segments  $e_i$  coincides with the boundary of  $K$ . Again, by known results [HS], [WS], the number of intersections along this envelope is  $O(n\alpha(n))$  and this bound is tight in the worst case.

The main characteristic of these special cases is that the complement of  $K$  is connected; equivalently, the boundary  $\partial(K)$  of  $K$  is a single connected curve. As will be apparent from the subsequent analysis, the crucial difficulty that we face in studying the general three-intersection problem is the presence of “holes” in  $K$  (as shown in Fig. 2). It should be noted that if we consider only the common boundary between  $K$  and a single connected component of the complement of  $K$ , then the number of intersections along that portion of  $\partial(K)$  is at most  $O(n\alpha(n))$ ; this is a variant of a more general result [GSS], [PSS], that the number of intersections along such a portion of  $\partial(K)$  is  $O(\lambda_{s+2}(n))$  if each  $\gamma_i$  is a Jordan arc and each pair of such arcs intersect in at most  $s$  points.

We deal with the presence of holes by lifting our arcs onto a surface  $\mathcal{R}$  on which the corresponding union of the interiors of these arcs is simply connected. This surface  $\mathcal{R}$  is a covering space for the region  $K$ . The holes in  $K$  arise because different “flaps” of the surface  $\mathcal{R}$  happen to overlap each other when projected canonically down to the plane.

The rest of the paper is organized as follows: Section 2 discusses the construction of the surface  $\mathcal{R}$  and a total ordering that results among the subarcs of the  $\gamma_i$  forming the boundary of  $K$ —incidentally, that section makes no use of the three-intersection property. In Section 3 this ordering is viewed as a symbol sequence in which five alternations of two symbols are forbidden but long repetitions are possible. Thus we nearly get a Davenport–Schinzel sequence of order 3, but not quite. Section 3 also gives an example showing that not only can  $\partial(K)$  have  $\Omega(n\alpha(n))$  intersections in the worst case (which can occur even if each  $\gamma_i$  is an extended segment, as discussed above) but that  $K$  can also have  $\Omega(n\alpha(n))$  holes. It is these holes that give rise to the long repetitions mentioned above. Finally, Section 4 deals with these uncontrolled repetitions using a global homotopy argument (here the three-intersection property is used strongly).

## 2. A Simply Connected Surface that Covers $K$

Let us consider the boundary  $\partial(K)$  of the region  $K$  that was defined earlier as the union of all regions  $K_i$  bounded by the arcs  $\gamma_i$ . The intersection points along the boundary  $\partial(K)$  subdivide it into a collection of primitive subarcs, each of which is a single connected piece of one of the original Jordan arcs  $\gamma_i$  for

some  $i$ ,  $1 \leq i \leq n$  (or of the  $x$ -axis, which by convention we count as one of our arcs). We call these primitive subarcs *segments*. Each endpoint of a segment is an intersection point. Thus the segments can naturally be grouped into connected components, each group bounding the same connected region of the complement of  $K$ . This complement consists of one infinite region, and zero or more bounded regions that we refer to as *holes*. The boundary of each hole is a circular sequence of segments, while the boundary of the infinite component is a bi-infinite chain of segments.<sup>2</sup>

In this section we derive the surprising fact that there is a natural linear ordering of all the segments in  $\partial(K)$ . This ordering stems from the existence of a (in general nonsimple) curve  $\Gamma$  that starts at  $-\infty$  on the  $x$ -axis and proceeds to the right, always following one of the  $\gamma_i$  consistently with its orientation, possibly switching arcs at intersection points but never retracing its steps, and finally ending at  $+\infty$  on the  $x$ -axis. We refer to each connected subarc that  $\Gamma$  traverses on a  $\gamma_i$  before switching to a different  $\gamma_j$  as an *edge* of  $\Gamma$ . The curve  $\Gamma$  traverses each segment in  $\partial(K)$  exactly once and satisfies two important further conditions:

- (a)  $\Gamma$  visits all segments coming from the same original arc  $\gamma_i$  in their natural order, and
- (b) all holes lie outside of  $\Gamma$ , in the sense that  $\Gamma$  can be continuously deformed over  $K$  so as to coincide with the  $x$ -axis.

This curve  $\Gamma$  is the projection onto the  $xy$ -plane of the boundary of the simply connected surface  $\mathcal{R}$  mentioned earlier. The surface  $\mathcal{R}$  itself can be obtained by gluing together copies of the  $K_i$  by a topological process described in detail below. As we mentioned earlier, this construction is independent of the three-intersection property (which is not assumed in this section). However, we still assume that any two distinct arcs  $\gamma_i$  and  $\gamma_j$  intersect only transversely and at a finite number of points. This allows us to describe the gluing process in finite (combinatorial) terms.

Consider the arrangement<sup>3</sup> defined on the  $xy$ -plane by the  $n$  bi-infinite arcs  $\gamma_i$ . This arrangement subdivides the plane into a number of connected two-dimensional cells that are joined together along pieces of our arcs. The finiteness assumption above implies that this arrangement has finite combinatorial complexity. We form a covering space for the cells in  $K$ , as follows. Let us take as many copies of each cell in the arrangement as there are arcs containing the cell in their interior. For instance, we take zero copies of each hole, one copy of each cell in  $K$  bordered by a segment not on the  $x$ -axis, and  $n$  copies of the lower halfplane. We call these copies *patches* to distinguish them from the underlying cells themselves. Initially we sew these patches together to form  $n$  surfaces<sup>4</sup>  $S_i$ , each surface corresponding to one of the  $K_i$ . We then repeatedly perform a “gluing” operation combining pairs of these surfaces into new surfaces, until

<sup>2</sup> By abuse of language, we also refer to the unbounded region as a hole.

<sup>3</sup> See [E] for more information on arrangements, in general.

<sup>4</sup> Technically, we use the word “surface” to refer to a 2-manifold with boundary.

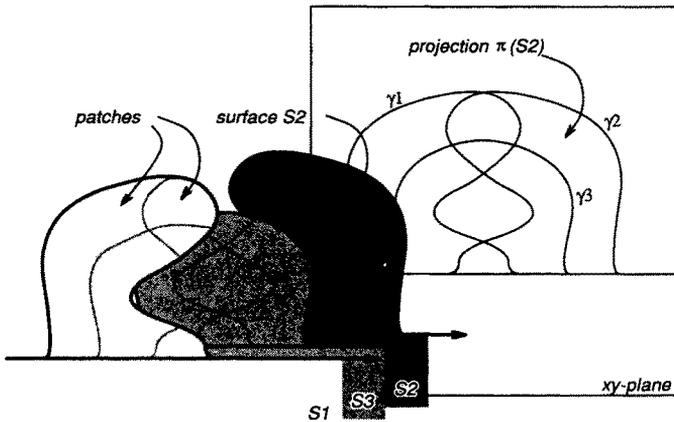


Fig. 4. An illustration of the surfaces  $S_i$ .

only one surface remains; this will be the desired surface  $\mathcal{R}$ . We now elaborate on this construction.

The initial surface  $S_i$ , for  $1 \leq i \leq n$ , is formed by taking a patch corresponding to each of the cells in  $K_i$  and then attaching these patches together along shared edges, by just imitating the way the corresponding cells attach to each other down in the  $xy$ -plane. It is clear that each  $S_i$  is a simply connected surface. Associated with each of the  $S_i$  is a canonical projection map  $\pi$  mapping points on  $S_i$  to the corresponding points in  $K_i$  on the  $xy$ -plane.<sup>5</sup> Figure 4 gives an example of three arcs  $\gamma_1, \gamma_2, \gamma_3$  and the corresponding surfaces  $S_1, S_2, S_3$ .

So far it may seem that we have done a lot of work for nothing: after all we only got the  $K_i$  back from the construction. But we now start a process of gluing the  $S_i$  together to form more complex surfaces. This process becomes confusing if we attempt to understand it while working on the  $xy$ -plane only. The reason is that some of the surfaces we form may contain several of the patches covering the same cell in the  $xy$ -plane. The proper way to keep track of which pairs of patches are to be identified (“glued”), and which are not, is to work with the more general concept of a surface (2-manifold with boundary).

The gluing process giving rise to these surfaces can be described at the top level as follows. Initially we put the  $n$  surfaces  $S_1, S_2, \dots, S_n$  in a list  $\mathcal{L}$ . Then we repeatedly perform the following operation until our list contains only a single surface: extract any two surfaces from  $\mathcal{L}$ , glue them together in a manner to be specified, and put the resulting surface back in  $\mathcal{L}$ . The surfaces that arise during the gluing process will be simply connected two-dimensional manifolds with boundary. Each such surface  $S$  will also have a canonical projection map  $\pi$  down to the  $xy$ -plane that, in general, will *not* be one-to-one. The surface  $S$  will naturally cover (be a covering space for) some subset of the  $K_i$ —those  $K_i$  corresponding to the  $S_i$ ’s that were glued together to form  $S$ . At any one time, for each  $K_i$  there will exist a unique surface  $S$  in our list  $\mathcal{L}$  in whose construction  $S_i$  was used.

<sup>5</sup> Our use of language suggests that we think of the  $S_i$  as lying “above” the  $xy$ -plane.

Furthermore, it will be possible to lift the region  $K_i$  homeomorphically and canonically onto the corresponding surface  $S$  via the inverse projection map  $\pi^{-1}$ .

During the gluing process certain pairs of points belonging to different surfaces, or even possibly to the same surface, become identified, or glued together. Two points can be identified only if they project to the same point on the  $xy$ -plane, according to their respective canonical projection maps. Actually, the gluing always works in a patch-by-patch basis. Two patches covering the same cell in the  $xy$ -plane are called *compatible*. Two compatible patches are either fully identified in the obvious way, or kept fully disjoint. Thus the gluing process can be described in combinatorial terms as acting on the finite collection of patches. Finally, a small point of terminology: a patch is called a *boundary patch* if one of its bounding edges is part of the boundary of the surface it lies on.

Let the contents of the list  $\mathcal{L}$  be  $T_1, T_2, \dots, T_k$  at some point during the gluing process. Let  $f$  denote the function such that  $f(i) = j$ ,  $1 \leq i \leq n$ ,  $1 \leq j \leq k$ , if the original surface  $S_i$  was used during the gluing process in forming  $T_j$ . We summarize below the invariants that we maintain:

- (a) each  $T_j$ ,  $1 \leq j \leq k$ , is a simply connected surface with boundary,
- (b) each  $T_j$ ,  $1 \leq j \leq k$ , singly covers the lower halfplane, and either does not cover, or singly covers, each cell adjacent to the lower halfplane,
- (c) for each  $T_j$ ,  $1 \leq j \leq k$ ,  $\pi(T_j) \subseteq K$ , where  $\pi$  is the canonical projection map of  $T_j$ ,
- (d) each boundary patch of a  $T_j$  is a boundary patch of an  $S_i$  for some  $i$  such that  $f(i) = j$ , and
- (e) each  $K_i$ ,  $1 \leq i \leq n$ , uniquely and homeomorphically lifts onto  $T_{f(i)}$  via  $\pi^{-1}$ .

We now give the details of the gluing process and prove that these invariants are maintained. Without loss of generality, assume that our next action on the list  $\mathcal{L}$  is to glue together  $T_1$  and  $T_2$ . We start by identifying the patches in  $T_1$  and  $T_2$  corresponding to the lower halfplane. We now put in a bag all the remaining patches of  $T_1$  and  $T_2$ , together with this single copy of the lower halfplane, and repeatedly perform the following operation. As long as there are two compatible patches in the bag that border<sup>6</sup> the same patch along the same edge, we identify them and put the single copy back in the bag. This process clearly terminates, as the number of patches originally in the bag is finite, and each identification reduces the number of patches by one. The remaining patches in the bag, when the gluing is finished, can be resewn together along shared seams by the original recipes provided by  $T_1$  and  $T_2$ . What is the nature of the resulting object  $T$ ? Before we address this question, let us work through this gluing process as applied to the three surfaces of Fig. 4.

Suppose  $S_1$  and  $S_2$  are chosen to be glued first. Besides the lower halfplane copies, only one more pair of patches needs to be identified, yielding the surface  $T$  shown in Fig. 5. Note that this surface has two overlapping flaps that were *not* glued together. Next we glue together  $T$  and  $S_3$ . Three pairs of patches border the  $x$ -axis and they get glued together to start with. Now the patches in  $T$

<sup>6</sup> We assume that gluing does not disturb any pre-existing adjacency relationships among patches.

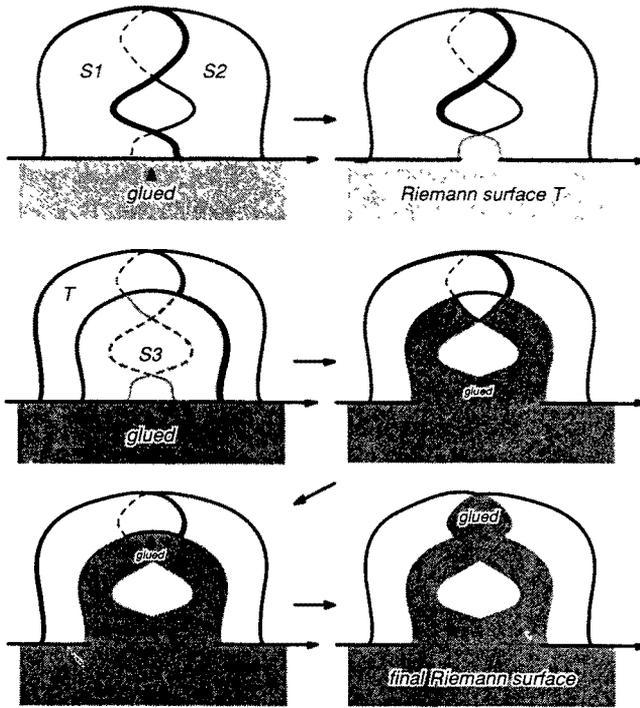


Fig. 5. The gluing process on the example of Fig. 4.

corresponding to the portions of the overlapping flaps covered by  $S_3$  are glued to the compatible patch of  $S_3$ , and therefore to each other. Finally, the patches in  $T$  corresponding to the portion of the overlapping flaps not covered by  $S_3$  are glued together, yielding the final surface  $\mathcal{R}$ . Note that this last gluing was between two patches of the same surface. Figure 5 illustrates this entire construction.

**Lemma 2.1.** *The result  $T$  of gluing together  $T_1$  and  $T_2$  as described above is a surface.*

*Proof.* Since  $T$  consists of surface patches sewn together along common boundaries, it can only fail to be a surface if it were to contain a “bifurcation”, as shown in Fig. 6. Such bifurcations can indeed be present in the middle of the

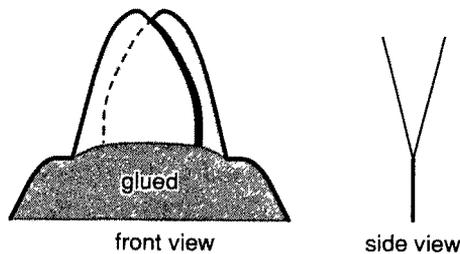


Fig. 6. A bifurcation.



gluing operation described above. But the presence of a bifurcation signifies that the gluing operation is not yet finished—there are more compatible patches that ought to be glued together. So when the gluing is finished, the remaining patches sewn together do form a surface.  $\square$

The surface  $T$  is clearly connected (everything connects to the patch covering the lower halfplane), but in fact more is true.

**Lemma 2.2.** *The surface  $T$  is simply connected.*

*Proof.* Simple connectivity for a surface means that any loop on it can be contracted to a point. We assume inductively that the surfaces  $T_1$  and  $T_2$  are simply connected. In our argument, we find it necessary to apply the concept of simple connectivity to the “surfaces with bifurcations” that arise as intermediate temporary objects during the gluing of  $T_1$  with  $T_2$  to produce  $T$ . Let us call these intermediate objects *semisurfaces*. For the sake of brevity we do not give formal definitions of these concepts here.

We start a gluing operation by identifying the patches of  $T_1$  and  $T_2$  corresponding to the lower halfplane. Any loop on this semisurface that does not go over the glued area is clearly a loop on either  $T_1$  or  $T_2$ , and therefore contractible. If a loop  $\gamma$  goes over the glued area and cannot be viewed as a loop lying entirely on  $T_1$ , or entirely on  $T_2$ , then  $\gamma$  can be written as  $\gamma = \alpha_1\alpha_2 \cdots \alpha_m$ , where each of the  $\alpha_i$  is an arc that can be drawn entirely on  $T_1$ , or entirely on  $T_2$ .<sup>7</sup> It is clear that the transition points between these arcs can be chosen to be points that, under the canonical mapping  $\pi$ , project on the  $x$ -axis. Let  $v_i$  be the transition vertex from  $\alpha_{i-1}$  to  $\alpha_i$  (we set  $v_0 = v_m$ ). Choose  $x$  to be any point in the glued area and connect  $x$  by paths  $\rho_i$  to each of the vertices  $v_i$ , where all these paths lie entirely in the glued region. Note that as a consequence each  $\rho_i$  corresponds to a path in  $T_1$ , as well as a path in  $T_2$ . Now  $\gamma$  is clearly homotopic to  $(\rho_1\alpha_1\rho_2^{-1})(\rho_2\alpha_2\rho_3^{-1}) \cdots (\rho_m\alpha_m\rho_1^{-1})$ . Each term in parentheses is a loop in either  $T_1$  or  $T_2$ , and therefore trivial. Thus  $\gamma$  itself is contractible.

We now prove that each further identification of two patches during our gluing operation preserves simple connectivity. Let  $\mathcal{S}$  denote the current simply connected semisurface and assume our next step is to glue together two patches  $Q$  and  $R$  to form a new patch  $U$ , and thus obtain a new semisurface  $\mathcal{S}'$ . By assumption,  $Q$  and  $R$  must be compatible and they must both border a common patch  $P$  of  $\mathcal{S}$ . Let  $\gamma$  be a loop in  $\mathcal{S}'$ . We argue in a manner similar to that of the previous paragraph. If  $\gamma$  is already a loop in  $\mathcal{S}$ , there is nothing to show. Otherwise, as above,  $\gamma$  can be written as  $\gamma = \alpha_1\alpha_2 \cdots \alpha_m$ , where each of the  $\alpha_i$  is an arc that can be drawn entirely on  $\mathcal{S}$ . Again, the breakpoints  $v_i$  between these arcs can be chosen to be points on the boundary of the patch  $U$ . Let  $x$  be a point on the common boundary between this patch and  $P$ . Connect  $x$  by a path  $\rho_i$  lying

<sup>7</sup> We assume here that the loop  $\gamma$  intersects the  $x$ -axis only finitely often, but our argument does not depend on this assumption. The same comment applies to general gluing step discussed in the sequel.

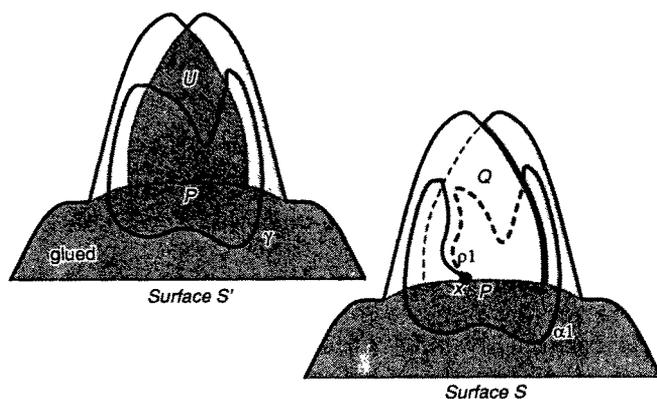


Fig. 7. The argument for simple connectivity.

entirely in  $U$  to each vertex  $v_i$ . The  $\rho_i$  thus correspond to paths in both  $Q$  and  $R$ . Therefore, if we write  $\gamma$  as  $(\rho_1\alpha_1\rho_2^{-1})(\rho_2\alpha_2\rho_3^{-1})\cdots(\rho_m\alpha_m\rho_1^{-1})$ , we see that it is equivalent to a loop on  $\mathcal{S}$ , and therefore contractible. Figure 7 illustrates this construction.

By continuing in this fashion, we can prove that the final semisurface of the gluing operation, which is the surface  $T$ , is also simply connected. For the reader familiar with topology, the argument given above is a relatively simple variation on the proof of van Kampen's theorem; see, for example, [M] or [Spa].  $\square$

The rest of the invariants are now easy to verify. If both  $T_1$  and  $T_2$  singly cover the same cell  $C$  bordering the lower halfplane, then clearly their (unique) patches covering  $C$  will be identified. Thus  $T$  also covers each cell bordering the lower halfplane either zero times, or one time. Since no patches are thrown away during the gluing process, it follows that boundary patches in  $T$  must come from boundary patches in either  $T_1$  or  $T_2$ . Let now  $S_i$  be one of the initial surfaces used in forming (say)  $T_1$ ; thus  $f(i) = 1$ . By induction, a homeomorphic copy of  $S_i$  exists on  $T_1$ , and therefore on  $T$ . The reason is that each of the patches in  $T_1$  that forms this copy of  $S_i$  survives also in  $T$  (though it may have been glued to other patches of  $T_1$  or  $T_2$ ). This copy of  $S_i$  in  $T$  is canonical in the following sense. Consider a cell  $C$  on the  $xy$ -plane that is part of  $K_i$ . Let  $\rho$  be a path connecting some point of this cell to the lower halfplane and lying entirely in  $K_i$ . The endpoint of  $\rho$  in the lower halfplane admits of a unique lifting onto  $T$ . By an obvious monodromy argument (see, for example, [Spr]) it follows that all of  $\rho$  admits of a unique lifting onto  $T$ . This uniquely identifies the patch in  $T$  covering  $C$  that is associated with the canonical lifting of  $K_i$  onto  $T$  via  $\pi^{-1}$ . These observations prove that invariant properties (a), (b), (d), and (e) hold; the truth of (c) is immediate.

As each gluing operation reduces the number of surfaces on the list  $\mathcal{L}$  by one, after  $n - 1$  such operations our list will contain a single surface. This is the desired final surface  $\mathcal{R}$ . We have derived the fact that  $\mathcal{R}$  is simply connected and that it exactly covers  $K$ , the union of the  $K_i$ 's (the latter follows because  $\mathcal{R}$  contains

copies of all the patches and consists of nothing but copies of such patches). Each  $K_i$  has a canonical lifting onto  $\mathcal{R}$  as a simply connected subsurface of  $\mathcal{R}$ . The desired curve  $\Gamma$ , whose existence was postulated at the beginning of this section, is simply the projection on the  $xy$ -plane of the boundary  $\partial(\mathcal{R})$  via the canonical projection map  $\pi$ .

It is clear that any segment in  $\partial(K)$  not on the  $x$ -axis is part of the border of some cell in  $K$  that corresponds to a unique patch in one of the original surfaces  $S_i$ . This patch cannot partake in any gluing and therefore survives as a boundary patch in  $\mathcal{R}$ . Thus it follows that  $\pi(\partial(\mathcal{R}))$  covers all the segments of  $\partial(K)$ . Now consider any one of the original arcs, say  $\gamma_i$ . Since a homeomorphic copy of  $S_i$  exists on  $\mathcal{R}$  as a simply connected subsurface, it follows that any portions of  $\partial(\mathcal{R})$  that come from lifting  $\gamma_i$  occur along  $\partial(\mathcal{R})$  in the same order as along  $\gamma_i$ . A rigorous version of this argument is given in [GSS]. Thus  $\Gamma$  visits segments from the same arc  $\gamma_i$  in an order consistent with their natural order along  $\gamma_i$ . Finally, note that the  $x$ -axis has a unique lift onto  $\mathcal{R}$  (it is, after all, one of the  $\gamma_i$ ) and  $\partial(\mathcal{R})$  can be homotoped over  $\mathcal{R}$  to this lift of the axis. Down on the  $xy$ -plane this homotopy becomes a homotopy of  $\Gamma$  over  $K$  to the  $x$ -axis. Thus all holes are outside  $\Gamma$ . We have now proved the desired theorem.

**Theorem 2.1.** *There exists a bi-infinite curve  $\Gamma$  that includes all of the segments bounding  $K$ , each exactly once. This curve traverses each segment in an order consistent with its own orientation and traverses different segments from the same arc in an order consistent with their ordering along the arc. Furthermore,  $\Gamma$  can be continuously deformed inside  $K$  so as to coincide with the  $x$ -axis.*

This curve  $\Gamma$  for the arrangement in Fig. 2 is shown in Fig. 8.

We close this section with a remark that will come in handy in the next section. Suppose that along the boundary of  $\mathcal{R}$  we have an edge coming from original arc  $\gamma_i$ , later one from  $\gamma_j$ , and then one from  $\gamma_i$  again (these edges need not be consecutive). Then we claim that to each of the two transitions corresponds a crossing of  $\gamma_i$  and  $\gamma_j$ . To see this, lift  $\gamma_i$  onto  $\mathcal{R}$  and cut the surface by following this lifted image from the first edge of  $\gamma_i$  along the boundary to the second. The arc  $\gamma_j$  also lifts onto the same surface and needs to get from the  $x$ -axis to the piece we have cut off in between the two edges of  $\gamma_i$ . This implies that there must be two crossings between the lifted images of the arcs on the surface and *a fortiori* on the plane.

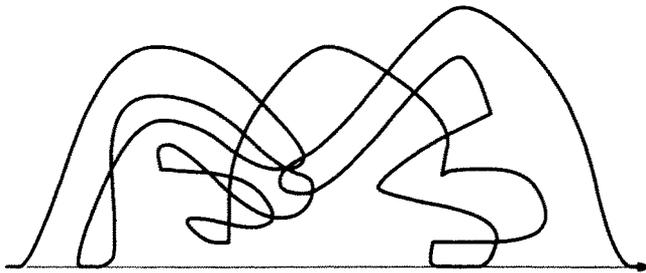


Fig. 8. The curve  $\Gamma$  corresponding to the arrangement of Fig. 2.

### 3. An Almost Davenport–Schinzel Sequence

We now concentrate our attention on the sequence  $\Sigma$  that consists of the segments in  $\partial(K)$  ordered according to the order implied by  $\Gamma$ . For the purposes of this section we think of  $\Sigma$  as a finite word over an alphabet whose symbols are the original arcs  $\gamma_i$ . Each segment  $s$  in  $\Sigma$  is given the symbol of the original arc that contains it. Here we start using the three-intersection property.

**Lemma 3.1.** *The sequence  $\Sigma$  does not contain an alternation among the original arcs of the form  $\gamma_i \cdots \gamma_j \cdots \gamma_i \cdots \gamma_j \cdots \gamma_i$ .*

*Proof.* Consider the boundary of the surface  $\mathcal{R}$ . An alternation of the specified kind along  $\Gamma$  clearly implies a similar alternation on the boundary of the surface. By the remark at the end of Section 2, each of the appearances of  $\gamma_j$  between two appearances of  $\gamma_i$  implies two crossings between the two arcs. These four crossings in total are distinct, as they occur along disjoint portions of  $\gamma_i$ , but this contradicts the three-intersection property.  $\square$

It is clear that along the boundary of  $\mathcal{R}$  the same labeling scheme produces an  $(n, 3)$  Davenport–Schinzel sequence  $\Sigma'$ , as any two consecutive edges must have distinct labels. By the results in [HS], the length of  $\Sigma'$  is bounded by  $O(n\alpha(n))$ . Unfortunately, down on the plane the same need not be true: along  $\Sigma$  adjacent equal elements can occur, as in the example of Fig. 9. In this case  $\Sigma = aaabbbccc$ . This happens because a single edge on the boundary of the surface  $\mathcal{R}$  can be broken into many pieces down in the plane by flaps of the surface crossing over it.

We also note that  $\Omega(n\alpha(n))$  is a lower bound on the length of  $\Sigma$  in the worst case, as well as on the number of holes. The bound on the length of  $\Sigma$  follows from the result of [WS] mentioned earlier. In [WS] Wiernik and Sharir construct a set of  $n$  line segments in the plane such that its upper envelope has  $\Omega(n\alpha(n))$  pieces. In their construction (in fact, in any arrangement of line segments) the upper envelope consists of  $O(n)$  convex bays, along each of which a particular segment can appear only once. We now use this example to construct a collection of  $O(n)$  three-intersecting Jordan arcs possessing many holes, as follows. Assume the  $x$ -axis lies below all segments. As described earlier, extend each segment to a trapezoid by dropping vertical lines from each of its endpoints to the  $x$ -axis. Now for each of the bays in the upper envelope we add a new simple Jordan

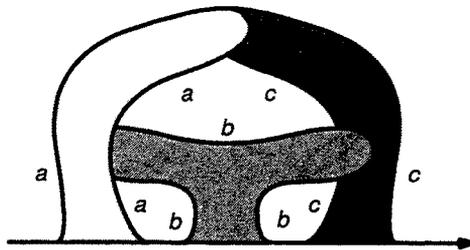


Fig. 9. Adjacent segments in  $\Sigma$  can come from the same edge.

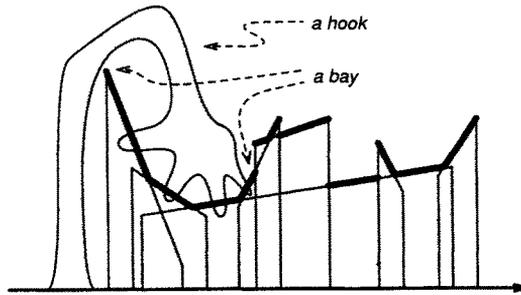


Fig. 10. The hooks used to construct an arrangement of three-intersecting Jordan arcs with many holes.

arc that starts and ends far to the left on the  $x$ -axis, but then comes in over the appropriate bay and cuts a piece off each of the segments in the bay, as exemplified in Fig. 10. We call these added Jordan arcs *hooks*. As Fig. 10 makes clear, each hook intersects a trapezoid in at most two points. Furthermore, the hooks can be drawn so they are disjoint among themselves. The result is an arrangement of  $O(n)$  three-intersecting simple Jordan arcs where the number of holes present is  $\Omega(n\alpha(n))$ ; there is one hole per inner vertex of each bay in the upper envelope. Note that this example gives us a construction where not only is the length of the contracted sequence  $\Sigma'$  at least  $\Omega(n\alpha(n))$ , but so is the number of repetitions in  $\Sigma$  (i.e., the quantity  $|\Sigma| - |\Sigma'|$ ).

Davenport-Schinzel theory in fact implies that we will get a nearly linear upper bound on the length of the contracted sequence  $\Sigma'$  as long as the maximum number of intersections between any pair of our Jordan arcs is bounded by any fixed constant  $s$ . The quadratic behavior in  $\Sigma$  for  $s \geq 4$  occurs exactly because of the breakage of single edges on the surface  $\mathcal{R}$  into many segments bearing the same label down on the plane. In the next section we show that under the three-intersection assumption such bad breakage cannot occur.

#### 4. Controlling the Fragmentation

Let us concentrate on a single edge  $e$  of  $\Gamma$ . We must understand how it can happen that  $e$  contains many segments; see Fig. 11 for an illustration. Let  $\alpha$  denote the original arc containing  $e$ , and let  $S_\alpha$  be the corresponding initial surface, a canonical homeomorphic copy of which lies in  $\mathcal{R}$ . Let  $e'$  denote the portion of  $e$  from the beginning of its first segment to the end of its last segment. The subarc  $e'$  consists of an alternation between segments and covered portions; these covered portions we term *windows*. As we remarked, windows can arise because different flaps of the surface  $\mathcal{R}$  can cross over  $e$ . The flaps are formed by the lifted interiors of other original arcs. In the subsequent analysis we make strong use of the three-intersection property.

Let  $\beta$  be an original arc that contributes to these flaps crossing over  $e'$  (so the homeomorphic copy of  $S_\beta$  in  $\mathcal{R}$  contains a portion that crosses over  $e$ ). Note

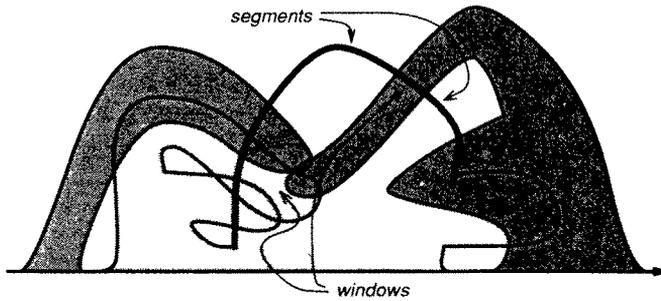


Fig. 11. The fragmentation of the darkened edge into segments and windows.

that along  $e'$ , the number of crossings of  $\alpha$  and  $\beta$  has to be even. The reason is that in effect  $e'$  starts and ends in holes, i.e., outside  $\beta$ . The same argument actually applies to the crossings along any window. Because of the three intersection property, it now follows that such an arc  $\beta$  must have exactly two intersections with  $e'$ , and these intersections must both occur in the same window. So there are only two ways in which  $\beta$  can cross over  $e'$ , as illustrated in Fig. 12: the arc  $\beta$  can either cross from the inside, or from the outside.

**Lemma 4.1.** *No original arc can cross over a window from the inside.*

*Proof.* Assume that  $\beta$  crosses over  $e'$  from the inside. Let  $p$  and  $q$  be the two crossings on  $e'$  where  $\beta$  escapes out of  $K_\alpha$ . Define  $\beta_p$  to be the subarc of  $\beta$  that connects  $p$  to the  $x$ -axis and does not contain  $q$ ; define  $\beta_q$  symmetrically. By the three-intersection property, at least one of the arcs  $\beta_p$  and  $\beta_q$  lies entirely inside  $K_\alpha$ . Thus there is a path from the lower halfplane to  $p$  (and  $q$ ) that lies inside both  $K_\alpha$  and  $K_\beta$ . By the definition of the gluing operation,  $S_\alpha$  must be glued to  $S_\beta$  along the portion of  $e$  between  $p$  and  $q$ , contradicting the assumption that  $e$  is a single edge of  $\Gamma$ .  $\square$

**Lemma 4.2.** *If the original arc  $\beta$  crosses the edge  $e$  belonging to original arc  $\alpha$  over a window from the outside, then the portion of  $\beta$  between its entry and exit into  $K_\alpha$  along that window is fully contained inside  $K_\alpha$ .*

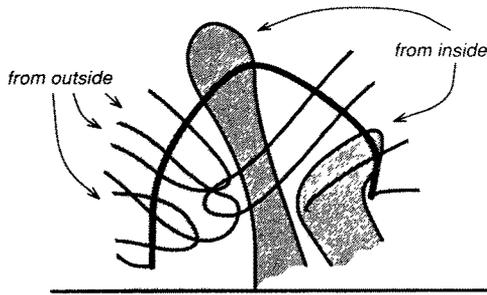


Fig. 12. The two ways an arc can cross over the darkened edge.

*Proof.* Obvious from the three-intersection property. □

Now that we have all this structure at our disposal we can proceed as follows. Order all the original arcs  $\gamma$ , in some arbitrary order. Choose the next unprocessed arc  $\alpha$  and consider all the other arcs crossing into  $\alpha$  from the outside, along windows occurring on edges of  $\alpha$ . Each intruding arc enters through only a single window, so we can partition the intruders according to the window through which they enter. Take all intruders entering through a particular window  $w$  and deform them as a group inside  $K_\alpha$  so as to bring them extremely close to the window  $w$ . By this we mean (1) that we can bring them so close to  $w$  that no arc except  $\alpha$  contains the deformed portions in its interior, and (2) that we maintain the combinatorial structure of the arrangement of these intruders among themselves inside  $K_\alpha$ . Notice that we can do this deformation so that no new intersections are introduced (though old intersections may be lost). This is so because there are no arcs crossing  $w$  from the inside, by Lemma 4.1. Also note that this deformation does not affect the covered region  $K$  and it maintains the three-intersection property. We do such a group deformation for every window  $w$  of  $\alpha$ . It is clear that the deformed portions do not interfere with each other, since windows are separated along  $\alpha$ . This completes the processing of  $\alpha$  and we now proceed similarly to the next original arc. Since all deformations associated with the processing of one arc happen entirely inside the interior of the arc and “clear out” that interior, subsequent deformations affect interiors of other arcs and hence do not interfere with the cleared region inside  $\alpha$ . Figure 13 illustrates this deformation process.

The end result of this deformation process is that for each window  $w$  of each arc  $\alpha$  we have cleared out a small corridor inside  $K_\alpha$  and near  $w$ . That corridor is covered by  $K_\alpha$  only. Furthermore, the corridors corresponding to different windows on the same edge can all be connected into a single connected region that is covered by  $K_\alpha$  only and includes on its boundary all segments on that edge. In Fig. 13, the corridor for one of the edges of  $\alpha$  is enclosed by a dashed line. For every edge  $e$  of  $\alpha$  let us now choose one point  $p_e$  arbitrarily in the combined corridor associated with  $e$ , and then connect  $p_e$  to the endpoints of the segments on the edge  $e$  by nonintersecting paths lying inside  $K_\alpha$  only. This is clearly possible. We do this for all of our original arcs. We also choose a point

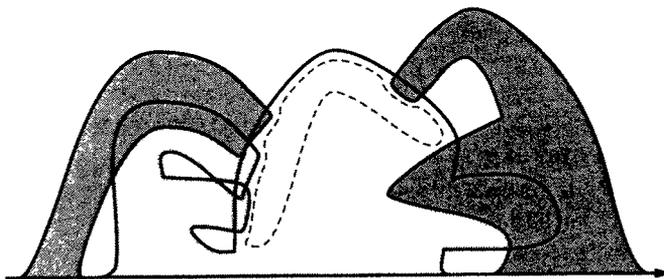


Fig. 13. The deformation process.

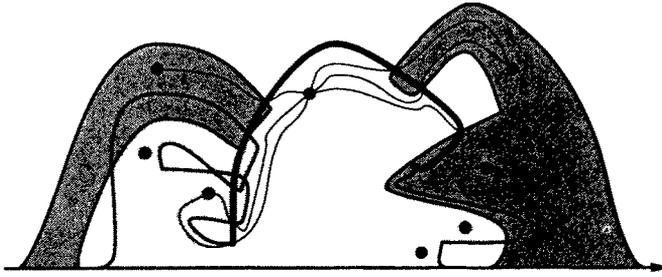


Fig. 14. The vertices of the graph  $G$  and the paths crossing the darkened edge.

in the lower halfplane and connect it by nonintersecting paths in the lower halfplane to all the segment endpoints that lie on the  $x$ -axis. Now consider the graph  $G$  whose vertices are these newly introduced points and whose edges are defined by the nonintersecting paths we just drew. Each graph edge corresponds to two of our paths coming together at a hole vertex. Figure 14 illustrates that construction. A similar construction was used in [KL\*].

Note the following facts. Firstly, the number of vertices of  $G$  is the number of edges of the surface  $\mathcal{R}$ , which in our case is  $O(n\alpha(n))$ . Secondly, every hole vertex gives rise to an edge of  $G$ , because it corresponds to an endpoint of two segments, one, say, along arc  $\alpha$ , and the other along arc  $\beta$ . Thirdly, the number of edges between two given vertices of  $G$  is at most three, by the three-intersection property. And fourthly, the way that we have drawn the edges of  $G$  shows that it is a planar graph, as each halfedge of it is covered by the interior of a single arc. From Euler's theorem we can finally conclude the desired result.

**Theorem 4.1.** *The maximum total number of sides in the exterior boundary of the union of  $n$  three-intersecting Jordan arcs is  $\Theta(n\alpha(n))$ .*

We remark that the same theorem holds if we arbitrarily choose for some of the  $\gamma_i$  the region  $K_i$  to be the one to the right of the arc (as we have consistently done up to now), and for others the one to the left of  $\gamma_i$  (or, equivalently, orient those  $\gamma_i$  the opposite way). The reason is that all of our arguments can be carried out on the sphere as well as on the plane. The  $\gamma_i$  there start and end on the equator, and each  $K_i$  includes the portion of the upper hemisphere to the right of the corresponding  $\gamma_i$ , as well as the lower hemisphere. If all the  $K_i$  together completely cover the sphere, then there is nothing to show. Otherwise we can carry out our surface construction for  $\mathcal{R}$  as before; we start by identifying the patches corresponding to the lower hemisphere and proceed to glue in the manner of Section 2. It is also easy to check that the deformation arguments needed in Section 4 are not affected by the couple of new cases that arise from doing the construction on the sphere. We end by repeating a remark we made in the Introduction: an interesting consequence of this extension of our result is that the exterior boundary of the union of  $n$  halfspaces bounded by three-intersecting bi-infinite curves has combinatorial complexity  $\Theta(n\alpha(n))$ .



## 5. Conclusions

We believe that the construction of the surface  $\mathcal{R}$  sheds some light on how the complexity of the exterior boundary of the union of a number of Jordan regions arises out of two different effects: the Davenport–Schinzel controlled behavior along the boundary of this surface and then the internal fragmentation introduced by the overlap of different flaps when projected down onto the plane.

We hope that this surface construction will also be found of use in algorithmic questions, since it provides a simply connected manifold of nearly linear boundary complexity (as long as the number of intersections per pair of arcs is bounded) that covers an area of the plane that can have quadratic complexity when explicitly represented.

## Acknowledgments

The authors would like to thank Sylvain Cappell, Edith Cohen, Lyle Ramshaw, Jorge Stolfi, and Bill Thurston for many useful discussions.

## References

- [A] Atallah, M.: Some dynamic computational geometry problems. *Comput. Math Appl.* **3** (1985), pp. 1171–1181.
- [E] Edelsbrunner, H.: *Algorithms in Combinatorial Geometry*. Springer-Verlag, New York (1987).
- [GSS] Guibas, L., Sharir, M., and Sifrony, S.: On the general motion-planning problem with two degrees of freedom. *Proceedings of the Fourth ACM Symposium on Computational Geometry* (1988).
- [HS] Hart, S., and Sharir, M.: Nonlinearity of Davenport–Schinzel sequences and generalized path compression schemes. *Combinatorica* **6** (1986), pp. 151–177.
- [KL\*] Kedem, K., Livne, R., Pach, J., and Sharir, M.: On the union of Jordan regions and collision-free translational motion amidst polyhedral obstacles. *Discrete Comput. Geom.* **1** (1986), pp. 59–71.
- [M] Massey, W. S.: *Algebraic Topology: An Introduction*. Springer-Verlag, New York (1977).
- [PSS] Pollack, R., Sharir, M., and Sifrony, S.: Separating two simple polygons by a sequence of translations. *Discrete Comput. Geom.* **3** (1988), pp. 123–136.
- [S] Sharir, M.: Almost linear upper bounds on the length of general Davenport–Schinzel sequences. *Combinatorica* **7** (1987), pp. 131–143.
- [Spa] Spanier, E.: *Algebraic Topology*. McGraw-Hill, New York (1966).
- [Spr] Springer, G.: *Introduction to Riemann Surfaces*. Chelsea, New York (1981).
- [WS] Wiernik, A., and Sharir, M.: Planar realization of nonlinear Davenport–Schinzel sequences by segments. *Discrete Comput. Geom.* **3** (1988), pp. 15–47.

Received July 19, 1988, and in revised form March 27, 1989.