

On the Union of Jordan Regions and Collision-Free Translational Motion Amidst Polygonal Obstacles*

Klara Kedem,¹ Ron Livne,¹ János Pach,³ and Micha Sharir^{1,2}

¹School of Mathematical Sciences, Tel Aviv University, Israel

²Courant Institute of Mathematical Sciences, New York University, New York, New York, USA

³Mathematical Institute of the Hungarian Academy of Sciences

Abstract. Let $\gamma_1, \dots, \gamma_m$ be m simple Jordan curves in the plane, and let K_1, \dots, K_m be their respective interior regions. It is shown that if each pair of curves $\gamma_i, \gamma_j, i \neq j$, intersect one another in at most two points, then the boundary of $K = \bigcup_{i=1}^m K_i$ contains at most $\max(2, 6m - 12)$ intersection points of the curves γ_i , and this bound cannot be improved. As a corollary, we obtain a similar upper bound for the number of points of local nonconvexity in the union of m Minkowski sums of planar convex sets. Following a basic approach suggested by Lozano Perez and Wesley, this can be applied to planning a collision-free translational motion of a convex polygon B amidst several (convex) polygonal obstacles A_1, \dots, A_m . Assuming that the number of corners of B is fixed, the algorithm presented here runs in time $O(n \log^2 n)$, where n is the total number of corners of the A_i 's.

1. Introduction

In this paper we consider the following restricted instance of the Piano Movers' problem [18]: Given a convex polygonal body B , free to translate (but not to rotate) in a 2-dimensional open region bounded by a collection of m convex polygonal obstacles A_1, \dots, A_m , and an initial and final configurations of B , we wish to determine whether there exists a (purely translational) continuous ob-

*Work on this paper by the second author has been supported in part by a grant from the Bat-Sheva Fund at Israel. Work by the fourth author has been supported in part by a grant from the U.S.-Israeli Binational Science Foundation.

stacle-avoiding motion of B between the two given configurations, and if so plan such a motion. In the more general problem considered in [18], the object B is also free to rotate. This leads to a more difficult problem, whose best solution to date runs in time $O(n^2 \log n)$ ([10], [20]; cf. also [14], [15]). Nevertheless, as noted in [11], in pragmatic applications it may be sufficient to consider only purely translational motions of B , or at most translational motions interleaved with one rotation, which is done in areas relatively “free” of obstacles. This simplified version of the motion planning problem has been considered by Lozano-Perez and Wesley in [11], and also in [16] (where B is assumed to be a circular disc, and where an $O(n \log n)$ motion planning algorithm is presented for this special case). The method presented by Lozano-Perez and Wesley [11] for the case of a polygonal object is given only in general terms, and no complexity analysis is provided. In this paper we follow the general scheme of [11], but develop it into an efficient algorithm which runs in time $O(n \log^2 n)$, where n is the number of obstacle corners. The algorithm is based on an interesting property of the union of certain Minkowski sums of convex 2-dimensional objects, which, in turn, is a simple consequence of the topological theorem for Jordan curves stated in the abstract and proved in Section 3. Some higher-dimensional generalizations of this result are discussed in Section 6. Section 2 introduces basic notation and terminology, and reviews the technique of Lozano-Perez and Wesley. Sections 4 and 5 present our efficient motion planning algorithm, and concluding remarks are given in Section 6.

2. The Approach of Lozano-Perez and Wesley to Translational Motion Planning

Before describing this approach, we begin with a few basic definitions.

Definition 2.1. The Minkowski (vector) difference of two planar sets A and B , denoted by $A - B$, is the set $\{p_1 - p_2 : p_1 \in A, p_2 \in B\}$.

Definition 2.2 (see [6]). The point p is a *local nonconvexity point* of a set S if each neighborhood of p contains two points $x, y \in S$ such that the segments $(px), (py) \subseteq S$ and the segment (xy) is not contained in S .

We will denote the interior of a set A by $int(A)$, and the boundary of A by $bd(A)$.

The motion planning problem studied in this paper can be stated as follows. Let B be a convex polygon in the plane, and let A_1, \dots, A_m be m closed convex polygonal obstacles, having disjoint interiors. B is free to translate in the plane, but must avoid collision with any of the obstacles. The observation in [11] is that we can replace this problem by that of planning a collision-free path for a single point between two specified positions amidst the “expanded obstacles” $K_i = A_i - B$, $i = 1, \dots, m$. (In these differences we use a standard placement of B , and it is also assumed that the origin falls on a point p of (this placement of) B). This modified problem is immediately solved, once we have computed the complement of the union of the expanded obstacles $K^c = (\cup_{i=1}^m K_i)^c$, which is the set FP of

free placements of the reference point p of B . Hence computing K is one of the main goals of this work.

It is easily seen that K is bounded by a collection of polygonal curves. Each convex corner of K is a convex corner of one of the expanded obstacles K_i , whereas the remaining corners of K are points of local nonconvexity, each of which is an intersection of the boundaries of two of the expanded obstacles.

To find the convex corners of K , we can use the well known fact (cf. [2], [7]) that the difference $A - B$ of two convex polygons A, B in the plane, having p, q corners respectively, is a convex polygon having at most $p + q$ corners. An algorithm that finds the corners of $A - B$ in time linear in $p + q$ is presented in [2].

Assume that B has k corners, and that the total number of corners of the obstacles A_1, \dots, A_m is $n = \sum_{i=1}^m n_i$. The total number of convex corners of K is $\sum_{i=1}^m (n_i + k) = n + mk$, so that if k is bounded by some fixed constant (as may well be the case in practice) then the number of convex corners of K is $O(n)$. In the following section we will show that the number of corners of local nonconvexity of K is only $O(m)$, where this property also holds for general convex sets A_i, B . This will therefore imply that the total number of corners of K is $O(n)$.

3. The Union of Planar Jordan Regions

In this section we derive an estimate on the number of points of local nonconvexity in the union of Minkowski differences of the sort considered above. This estimate will imply that the complexity of such a union is not too large, and that it can be calculated rather efficiently.

Theorem 3.1. *Let A_1, \dots, A_n, B be closed convex sets in the plane such that A_1, \dots, A_n have disjoint interiors, and let $K_i = A_i - B$ for $i = 1, \dots, n$. Then $K = \bigcup_{i=1}^n K_i$ has at most $\max(2, 6n - 12)$ points of local nonconvexity.*

Theorem 3.1 is, in fact, an easy consequence of a general topological result (Theorem 3.2) on families of Jordan curves. To formulate this, we need some preparation.

A collection $\Gamma = \{\gamma_i\}_{i=1}^n$ of simple closed Jordan curves in the plane is *admissible* if for each $i \neq j$ the intersection $\gamma_i \cap \gamma_j$ consists of two crossing points or is empty.

Remark. A similar notion as defined in [5, p. 55] is that of an *arrangement* of curves in the plane. There one requires *every* pair of curves γ_i, γ_j in the arrangement to intersect in two crossing points. Our admissible collections are also mentioned in [5, p. 68], and are called there *weak arrangements of curves*, but the issues that concern us here are not discussed in [5].

For a simple closed Jordan curve γ let $K(\gamma)$ denote the closure of the interior region bounded by γ ; thus $\gamma = bdK(\gamma)$. For a collection $\Gamma = \{\gamma_i\}_{i=1}^n$ of such curves let $K(\Gamma) = \bigcup_{i=1}^n K(\gamma_i)$. Suppose Γ is admissible, and let $I(\Gamma) = \bigcup_{i \neq j} (\gamma_i \cap \gamma_j)$, and $E(\Gamma) = I(\Gamma) \cap bdK(\Gamma)$; thus $I(\Gamma)$ is the set of all intersection points of the curves in Γ , whereas $E(\Gamma)$ is the set of all such intersection points which lie

on the boundary of $K(\Gamma)$. Trivially the cardinality $\#I(\Gamma) \leq n(n-1)$, and this bound is easily achieved for any n . However, we will prove

Theorem 3.2. *Suppose $\Gamma = \{\gamma_i\}_{i=1}^n$ is an admissible collection of simple closed Jordan curves. If $n \geq 3$ then $\#E(\Gamma) \leq 6n - 12$. Moreover for each $n \geq 3$ there exist admissible collections Γ_n of n curves for which $\#E(\Gamma_n) = 6n - 12$.*

Remark. It is not hard to modify the proof so that the Theorem also holds for any collection $\Gamma = \{\gamma_i\}_{i=1}^n$ of simple closed Jordan curves for which $\#(\gamma_i \cap \gamma_j) \leq 2$ for $i \neq j$. We will refer to such collections as being *weakly admissible*.

Proof of Theorem 3.2. Assume without loss of generality that our curves are in *general position*, i.e. no three of them have a point in common. (Otherwise we can slightly deform some of the γ_i 's, without decreasing $\#E(\Gamma)$, so as to satisfy this condition.) A curve $\gamma_i \in \Gamma$ is called *redundant* if $\gamma_i \subset \bigcup_{j \neq i} K(\gamma_j)$. Let

$$\begin{aligned} r_1(\Gamma) &= \# \{ \text{redundant curves in } \Gamma \} \\ r_2(\Gamma) &= \# \{ (i, j) \mid i \neq j, \gamma_i \cap \gamma_j \neq \emptyset \} \\ r_3(\Gamma) &= \# \{ (i, j, k) \mid i, j, k \text{ are distinct, } K(\gamma_i) \cap K(\gamma_j) \cap K(\gamma_k) \neq \emptyset \}. \end{aligned}$$

The proof will proceed by induction on the quadruples $\mathbf{r}(\Gamma) = [r_i(\Gamma)]_{i=1}^3$ in lexicographical order. The basis for the induction is provided by Proposition 3.3 below, which assumes that only $r_2(\Gamma)$ is nonzero. Each of the remaining induction steps in the proof of the theorem proceeds by taking an admissible collection Γ and deforming or eliminating some of its curves to obtain a new admissible collection Γ' having a simpler structure (that is $\mathbf{r}(\Gamma') \prec \mathbf{r}(\Gamma)$) but such that $\#E(\Gamma') \geq \#E(\Gamma)$, so that the induction hypothesis implies the desired inequality for Γ too.

Proposition 3.3. *Suppose $\Gamma = \{\gamma_i\}_{i=1}^n$ is an admissible collection of curves in general position, and $r_1(\Gamma) = r_3(\Gamma) = 0$. Then $\#E(\Gamma) \leq 6n - 12$.*

Proof. First note that in this special case we have $E(\Gamma) = I(\Gamma)$. For each $i \neq j$ the set $I_{ij} = \gamma_i \cap K(\gamma_j)$ is a closed, connected (possibly empty) arc. For each i let

$$D_i = K(\gamma_i) - \bigcup_{j \neq i} (\text{int } \gamma_j).$$

Then D_i is nonempty and bdD_i is a simple closed Jordan curve, and in fact

$$bdD_i = \gamma_i - \left(\bigcup_{j \neq i} I_{ij} \right) \cup \left(\bigcup_{j \neq i} I_{ji} \right).$$

For each i choose an arbitrary point $p_i \in \text{int } D_i$, and for each $j \neq i$, if $\gamma_j \cap \gamma_i \neq \emptyset$ choose an arbitrary point $q_{ij} \in I_{ji}$ which is not an endpoint. Then we can connect p_i to all such q_{ij} by simple arcs α_{ij} which are pairwise non-intersecting except at




Fig. 1

p_i , and contained in $int D_i$ except for their endpoints q_{ij} . Likewise we can connect q_{ij} to q_{ji} by a simple arc $\beta_{ij} = \beta_{ji}$ contained in the interior of the simple closed Jordan curve $I_{ij} \cup I_{ji}$ (this interior being also equal to $int \gamma_i \cap int \gamma_j$). See Fig. 1 for an illustration of these concepts.

We obtain this way a graph $G(\Gamma)$ whose vertices are the points p_i such that for each $i \neq j$ for which $\gamma_i \cap \gamma_j \neq \emptyset$, $G(\Gamma)$ contains the edge $\alpha_{ij} \cup \beta_{ij} \cup \alpha_{ji}$ connecting p_i and p_j . The above discussion plainly implies that $G(\Gamma)$ is a planar graph. Hence

$$e(G(\Gamma)) \leq 3v(G(\Gamma)) - 6,$$

where $e(G(\Gamma))$ and $v(G(\Gamma))$ denote the number of edges and the number of vertices of $G(\Gamma)$, respectively. But $v(G(\Gamma)) = n$ and $e(G(\Gamma)) = \frac{1}{2} \#E(\Gamma)$, which implies the proposition. \square

Assume now that Γ is any admissible collection of curves in general position, and that Theorem 3.2 is true for all collections with the same properties, preceding Γ lexicographically.

Step 1. If $r_1(\Gamma) > 0$, then choose a redundant curve γ_i in Γ , and replace Γ by $\Gamma' = \Gamma - \{\gamma_i\}$. Clearly Γ' remains admissible and in general position, $\#E(\Gamma') \geq \#E(\Gamma)$, while $r_1(\Gamma') < r_1(\Gamma)$, so that the induction step in this case is completed.

Step 2. Suppose that $\Gamma = \{\gamma_i\}_{i=1}^n$ is admissible, in general position, $r_1(\Gamma) = 0$, and $r_3(\Gamma) > 0$. Without loss of generality we may assume that $K(\gamma_1) \cap K(\gamma_2) \cap K(\gamma_3) \neq \emptyset$, and that for some $k \geq 3$ we have $\gamma_1 \cap K(\gamma_i) = I_i \neq \emptyset$ for $i \leq k$, and $\gamma_1 \cap \gamma_i = \emptyset$ for $i > k$. The following lemma is trivial (use e.g. Schoenflies' Theorem, cf. [12]).

Lemma. Let $A = \{z \in C \mid \frac{1}{2} \leq |z| \leq 2\}$. After applying a homeomorphism of the plane onto itself we may assume

1. γ_1 is the unit circle S^1 .
2. For $i > k$, $K(\gamma_i) \cap A = \emptyset$.
3. For $i = 2, \dots, k$ we have $K(\gamma_i) \cap A = \{z \in A \mid a_i \leq arg z \leq b_i\}$ for appropriate directions a_i, b_i . (See Fig. 2.)

Suppose now that Γ is as in the Lemma, so that $\{I_i\}_{i=2}^k$ are angular intervals on S^1 .




Fig. 2

Subcase A. If for some $i = 2, \dots, k$ the interval I_i is a subset of $\bigcup_{j=2, j \neq i}^k I_j$, replace γ_i by

$$\gamma_i^* = \gamma_i - \{z \in C \mid |z| \leq \frac{3}{2}\} \cup \{z \in C \mid |z| = \frac{3}{2}\}.$$

(See Fig. 3.)

The resulting collection is still in general position, $r_1 = 0$, but r_2 is smaller. Indeed $\gamma_i^* \cap \gamma_1 = \emptyset$ whereas $\gamma_i \cap \gamma_1 \neq \emptyset$, and no new intersections are added: if γ_i^* intersects some γ_j for $j \neq 1, i$ so did γ_i . Furthermore $\#E(\Gamma)$ has not changed, since the two eliminated intersections did not belong to $E(\Gamma)$. Thus the induction step can now be completed in this case.

Subcase B. Suppose finally that Γ is as above, but that for each $i = 2, \dots, k$ the interval I_i is not contained in $\bigcup_{j=2, j \neq i}^k I_j$. Since $r_1(\Gamma) = 0$ we may assume that $1 \in S^1 - \bigcup_{i=1}^k K(\gamma_i)$. Let $I_i = \{z \in S^1 \mid a_i \leq \arg z \leq b_i\}$. Since Γ is in general position, the a_i 's and b_i 's are all distinct. Also, since we now assume that no I_i is contained in $\bigcup_{j=2, j \neq i}^k I_j$, we may assume that $a_2 < a_3 < b_2 < b_3$, and also that $0 < a_2 < a_3 < \dots < a_k < 2\pi$. But then we can replace γ_2 by

$$\begin{aligned} \gamma_2^* = & \gamma_2 - \{z \mid |z| \leq \frac{3}{4}\} - \{z \mid |z| \leq \frac{3}{2}, \arg z = b_2\} \\ & \cup \left\{z \mid |z| = \frac{3}{4}, a_2 \leq \arg z \leq \frac{a_2 + a_3}{2}\right\} \\ & \cup \left\{z \mid \frac{3}{4} \leq |z| \leq \frac{3}{2}, \arg z = \frac{a_2 + a_3}{2}\right\} \\ & \cup \left\{z \mid |z| = \frac{3}{2}, \frac{a_2 + a_3}{2} \leq \arg z \leq b_2\right\}. \end{aligned}$$




Fig. 3




Fig. 4

In other words, we pull γ_2 and γ_3 somewhat apart, so as to move their intersection point within $K(\gamma_1)$ to the exterior of that region (see Fig. 4).

It is easily seen that this deformation of γ_2 keeps $r_1(\Gamma)$ equal to zero, and does not increase $r_2(\Gamma)$. Furthermore, $r_3(\Gamma)$ has decreased, because now $K(\gamma_1) \cap K(\gamma_2) \cap K(\gamma_3) = \emptyset$, whereas $\#E(\Gamma)$ has not decreased. Thus the induction step is completed in this case too, thereby completing the proof of the first part of Theorem 3.2.

To show that the bound obtained is actually tight in the worst case, we give an example in which the curves in Γ are all circles (Γ is always (weakly) admissible in this case). We draw $n \geq 3$ circles in the plane as follows. The first three circles are drawn so that each pair of them intersect, and such that their union is homeomorphic to an annulus, i.e. has a hole in the center. Then, inductively, we draw each additional circle inside a hole bounded by three arcs of previously drawn circles so that it intersects each of those arcs at two points, thereby leaving three smaller holes out of the original hole (see Fig. 5).

Thus the first three circles contribute six points to $E(\Gamma)$ and each additional circle contributes six more points, yielding altogether $6n - 12$ points in $E(\Gamma)$. This completes the proof of the theorem. \square

Remarks. (1) If we relax the requirement that each pair of the curves in Γ intersect in at most two points, the size of $E(\Gamma)$ can increase significantly. For example, if each pair of curves in Γ is allowed to intersect as much as a maximum of four times, $\#E(\Gamma)$ can become $\Omega(n^2)$.




Fig. 5

(2) Theorem 3.2 can be generalized to the case of Jordan curves on the unit sphere S^2 , where each curve γ_i partitions S^2 into an “interior” region and an “exterior” one. Indeed, if the union K of the closures $K(\gamma_i)$ of the interiors of the curves γ_i is the whole of S^2 then there is nothing to prove. Otherwise choose a point p in $S^2 - K$ and map $S^2 - \{p\}$ onto the plane, thereby reducing the situation to that assumed in Theorem 3.2.

(3) As already noted, if each curve γ_i is a circle then clearly $\Gamma = \{\gamma_i\}_{i=1}^n$ is admissible. In this special case however the property established in Theorem 3.2 is already known; it follows for example from the properties of Voronoi diagrams for a set of intersecting discs in the plane (cf. [8], [19], [23]). Another proof was given by Pach (cf. [21]).

(4) A result related to Theorem 3.2, concerning the maximal number of “osculation points” of admissible collections of Jordan curves appears in [4].

Proof of Theorem 3.1. By an easy compactness argument, it suffices to prove the assertion in the special case where (i) the sets A_1, \dots, A_n are disjoint, and (ii) each of A_1, \dots, A_n and B has a smooth, strictly convex boundary. However, in this case we have

Proposition 3.4. *For every pair i, j ($1 \leq i \neq j \leq n$), $bd(A_i - B)$ and $bd(A_j - B)$ have at most two points in common.*

Proof. Assume the contrary, and let $\epsilon_0 = \sup \epsilon$, where the supremum is taken over all $\epsilon > 0$ for which $|bd(A_i - \epsilon B) \cap bd(A_j - \epsilon B)| \leq 2$ for all $\epsilon' \leq \epsilon$.

If there exists a point $p \in bd(A_i - \epsilon_0 B) \cap bd(A_j - \epsilon_0 B)$ such that $A_i - \epsilon_0 B$ and $A_j - \epsilon_0 B$ have a common supporting halfplane at p , then it is easy to show that $A_i \cap A_j \neq \emptyset$, which contradicts our assumptions. If such a point does not exist, then it is easy to check that one can find a sufficiently small positive δ with the property that $|bd(A_i - \epsilon B) \cap bd(A_j - \epsilon B)|$ is constant on the interval $(\epsilon_0 - \delta, \epsilon_0 + \delta)$, contradicting the definition of ϵ_0 . Details are left to the reader. \square

Proposition 3.4 implies that $\{bdK_i\}_{i=1}^n$ is a (weakly) admissible collection of closed Jordan curves. (As a matter of fact, slightly perturbing B we can also assume that $\{bdK_i\}_{i=1}^n$ is admissible.) Since every point of local nonconvexity of $K = \bigcup_{i=1}^n K_i$ is an intersection of two different bdK_i 's, Theorem 3.1 now follows immediately from Theorem 3.2. \square

Remark. Going back to the original problem of planning a purely translational motion of a convex object B amidst convex interior-disjoint obstacles A_1, \dots, A_n , Theorem 3.1 can now be interpreted as stating that there exist at most $6n - 12$ positions of B in which it touches simultaneously two distinct obstacles, assuming that these obstacles are in *general position*, meaning that no two sides of the expanded obstacles K_i overlap.

4. Efficient Calculation of K

In this section we describe an efficient algorithm which calculates the contour (i.e. boundary) of $K = \bigcup_{i=1}^n K_i$ (see Section 2). The algorithm is based on a method

due to Ottman, Widmeyer and Wood [17] for calculating the boundary of the union of several superimposed polygonal planar regions, which is related to a technique due to Bentley and Ottmann [3] for counting and reporting intersections in a collection of planar line segments. These techniques run in time $O((n + t)\log n)$, where n is the number of line segments, and t is the number of intersection points between these segments. Since naive application of these techniques to the expanded polygons bdK_1, \dots, bdK_m may encounter quadratically many intersections of these sets, we therefore combine this technique with the following divide-and-conquer approach.

Algorithm (Divide and conquer algorithm for the calculation of $\cup K_i$.)

1. Calculate all the K_i 's.
2. Recursively find $G = \cup_{i \in g} K_i$ and $H = \cup_{i \in h} K_i$, where

$$g = \left\{ 1, \dots, \left\lfloor \frac{m}{2} \right\rfloor \right\}, h = \left\{ \left\lfloor \frac{m}{2} \right\rfloor + 1, \dots, m \right\}.$$

3. Find the contour of $K = G \cup H$, using the Ottmann-Widmeyer-Wood approach.

Following Theorem 3.1, step 3 of the algorithm runs in time $O(n \log n)$, since it implies that the number of corners in G and H , as well as the number t of intersection points between them, is $O(n)$. Hence the algorithm's time complexity is $O(n \log^2 n)$.

Remark. The above algorithm can be generalized in a straightforward manner, to yield an efficient algorithm for the calculation of the union $K(\Gamma)$ for an admissible collection Γ of Jordan curves, assuming that each of the curves γ_i in Γ has a relatively simple shape.

5. The Connected Regions of *FP*

In this section we sketch an algorithm that provides us with a convenient representation of the connected components of *FP*, using an "inclusion tree" which can be constructed in time $O(n \log n)$. $K = \cup K_i$ is a general, possibly not connected, planar region with polygonal boundaries. Since each edge of K is a border line between an expanded obstacle and *FP*, it is more convenient to describe the boundary of K as a set of simple polygons of two types: *free* simple polygons L , whose interior, in a sufficiently small neighborhood of L , is contained in *FP*, whereas their exterior, in a similarly sufficiently small neighborhood, is contained in the expanded obstacles, and *non-free* simple polygons that border expanded obstacles on their interior side and *FP* on their exterior side (see Fig. 6; in non-general positions of A_1, \dots, A_n, B the structure of K may become more degenerate, but still manageable).

The root of the inclusion tree T is the outermost boundary of K , if K^c is bounded, or the entire plane otherwise, and is always free. A simple polygon L_1 is a direct son of another polygon L_2 if L_1 is immediately contained in L_2 , i.e. there does not exist another simple boundary polygon contained in L_2 that includes L_1 .




Fig. 6

Each connected component C of FP can then be represented by a free node L of T and its direct (non-free) children L_1, \dots, L_p . Here L is the exterior boundary of C and L_1, \dots, L_p are the connected components of its interior boundary. This representation of K is easy to construct in time $O(n \log n)$ using a straightforward sweeping technique.

Specifically, we sort all the corners of the boundary of FP in ascending x order, and process them from left to right, updating the list of intersections of FP with a vertical scan-line as it sweeps through each of these corners, in a manner similar to that described in [13]. Whenever the sweeping process encounters a leftmost corner of some component L of the boundary of FP , it locates the two boundary components L_1, L_2 lying immediately above and below L . Suppose L is non-free, if either L_1 or L_2 is free, then it must be the father of L in T ; otherwise L_1 and L_2 are both brothers of L in T ; this enables us to insert L properly into T . Similar and symmetric rules govern the handling of a free component L .

Having this representation available, we can easily solve the original translational motion planning problem as follows: given two placements Z_1, Z_2 of (the reference point p on) B in FP , draw the segment Z_1Z_2 , and find the boundary component L_1 (respectively L_2) intersecting Z_1Z_2 nearest to Z_1 (respectively Z_2). (If no such component exist, obviously B can be moved to Z_2 from Z_1 along Z_1Z_2). Then locate L_1, L_2 in the inclusion tree T , and determine from their relationship in T whether they bound the same connected component C of FP . If not, no motion between Z_1 and Z_2 is possible. Otherwise a canonical (but non-optimal) motion of B from Z_1 to Z_2 can be constructed in linear time by moving along Z_1Z_2 , and by going around each boundary component of C which intersects Z_1Z_2 from one such intersection to the next one; (see Fig. 7).




Fig. 7

(The step that just decides whether Z_1 and Z_2 lie in the same connected component of FP can be performed by a faster ($O(\log n)$) point location algorithm (as in [9]), after an alternative $O(n \log n)$ preprocessing of K).

If one seeks optimal (shortest) motion from Z_1 to Z_2 , one then faces the problem of computing shortest paths in 2-dimensional polygonal spaces, which can be done in time $O(q^2) = O(n^2)$, where q is the number of corners in the connected component C of FP containing Z_1, Z_2 (see [1], [22]).

6. Generalizations, Concluding Remarks

Let $\Sigma = \{\sigma_i\}_{i=1}^n$ be a collection of two dimensional surfaces in E^3 which satisfy the following properties:

- (i) Each σ_i is homeomorphic to the unit sphere S^2 .
- (ii) For each $i \neq j$, $\sigma_i \cap \sigma_j$ is either homeomorphic to S^1 or is empty.
- (iii) For each triple i, j, k of distinct indices, $\sigma_i \cap \sigma_j \cap \sigma_k$ is either homeomorphic to S^0 or is empty.

For example, if each of the σ_i is a sphere then conditions (i)–(iii) hold, assuming non-degeneracy of the configuration of these spheres (that is, assuming that no pair of spheres is tangent to one another and that no sphere is tangent to the intersection curve of two other spheres; nevertheless, as in the 2-dimensional case, Theorem 6.1 below will also hold if such a degeneracy occurs).

Let $K(\sigma_i)$ be the closure of the interior region bounded by σ_i , and let $K(\Sigma) = \bigcup_{i=1}^n K(\sigma_i)$. Denote by $I(\Sigma)$ the set of all points of triple intersection of the surfaces in Σ , and let $E(\Sigma) = I(\Sigma) \cap bdK(\Sigma)$. We have trivially $\#I(\Sigma) \leq n(n-1)(n-2)/3$ and this bound can be easily achieved for an appropriate collection of n spheres. However we have

Theorem 6.1. *If $\Sigma = \{\sigma_i\}_{i=1}^n$, $n \geq 4$, is a collection of two dimensional surfaces in 3-space satisfying conditions (i)–(iii) listed above, then $\#E(\Sigma) \leq 2n(n-3)$.*

Proof. Fix one of the surfaces σ_i and consider the collection of curves $\gamma_j = \sigma_i \cap \sigma_j$, $j \neq i$, drawn on it. Some of these curves can be empty, but each nonempty γ_j is a closed Jordan curve on σ_i , and for each $j \neq k \neq i$ we have $\#(\gamma_j \cap \gamma_k) \leq 2$. Hence we can define $K(\gamma_j) = \sigma_i \cap K(\sigma_j)$, for $j \neq i$, and apply Theorem 3.2 (or rather its generalization as in Remark (2) in Section 3) to the collection $\Gamma = \{\gamma_j\}_{j \neq i}$, to obtain $\#(E(\Sigma) \cap \sigma_i) \leq 6(n-1) - 12 = 6(n-3)$. Repeating this argument for each of the surfaces σ_i , and observing that each point in $E(\Sigma)$ will be counted that way three times, we conclude that $\#E(\Sigma) \leq n/3 \cdot 6(n-3) = 2n(n-3)$. \square

Corollary 6.2. *Given n distinct balls B_1, \dots, B_n in 3-space, no pair of which are tangent to one another, the number of triple intersections of these balls which lie on the boundary of their union is at most $2n(n-3)$.*

Remarks. (1) We do not know whether this bound is tight. However, even when the surfaces in Σ are all spheres, $E(\Sigma)$ can contain $\Omega(n^2)$ points, so that in the worst case the above bound is tight up to a multiplicative constant. To see such

an example, let $n = 2k$; let $\sigma_1, \dots, \sigma_k$ be k spheres whose centers are the points $[\cos(2\pi i/k), \sin(2\pi i/k), 0]$, $i = 1, \dots, k$, and whose radii are all equal to some r_0 lying strictly between $\sin(\pi/k)$ and 1, so that the boundary of the union $K_0 = \bigcup_{i=1}^k K(\sigma_i)$ is homeomorphic to a torus. Note that the intersection curves of these first k spheres, which lie on the boundary of K_0 , are k circles $\gamma_1, \dots, \gamma_k$ of equal radii, such that the planes containing them all pass through the z -axis, and such that all these circles can be obtained by a revolution of one of them about the z -axis. The remaining k spheres are arranged so that their centers all lie on the z -axis, and such that each of them intersects each of the curves γ_i at a pair of points. By choosing the centers and the radii of these k additional spheres in an appropriate manner, we can ensure that each of these intersection points belongs to $E(\Sigma)$, so that $\#E(\Sigma) = 2k^2 = n^2/2$.

(2) Theorem 6.1 can be generalized to surfaces of arbitrary dimension, in a straightforward way. In particular, this generalization will imply that for any collection of n $(d-1)$ -spheres in E^d , the total number of points of intersection of d of these spheres which lie on the boundary of their union is $O(n^{d-1})$.

(3) The motion planning algorithm described in Section 4 has a straightforward generalization to the case in which the obstacles A_i can be non-convex. This follows from the fact that any non-convex polygonal obstacle is the union of openly disjoint convex sets, and Theorem 3.1 is valid for collections like that, too.

(4) If the moving body B is non-convex, then FP can have $\Omega(n^2)$ connected components. For example, take B to be an object consisting of two orthogonal segments meeting at a common endpoint, and let the obstacles be arranged as in Fig. 8. (We are indebted to S. Sifrony for this observation).

(5) The generalization of the problems studied in this paper to three dimensions leads to the problem of planning collision-free translational motion of a convex polyhedral body among a collection of convex polyhedral obstacles in 3-space. The problem of obtaining sharp bounds on the number of corners of the union of the corresponding expanded obstacles is still open.

References

1. T. Asano, T. Asano, L. Guibas, J. Hershberger, and H. Imai, Visibility polygon search and Euclidean shortest paths, Proc. 26th IEEE Symp. on Foundations of Computer Science, 1985, pp. 155–164.
2. R. V. Benson, Euclidean Geometry and Convexity, McGraw-Hill, 1966, pp. 97–113.
3. J. L. Bentley and A. Ottmann, Algorithms for reporting and counting geometric intersections, IEEE Trans. on Computers, Vol. C-28 (1979), pp. 643–647.




Fig. 8

4. P. Erdős and B. Grünbaum, Osculation vertices in arrangements of curves, *Geometriae Dedicata* 1 (1973) 322–333.
5. B. Grünbaum, Arrangements and Spreads, Regional Conference Series in Mathematics, Vol. 10, Conference Board of the Mathematical Sciences, American Mathematical Society, Providence R.I. 1972.
6. M. D. Guay and D. C. Kay, On sets having finitely many points of local nonconvexity, *Israel J. Math.* 8 (1970), 39–52.
7. L. Guibas, L. Ramshaw, and J. Stolfi, A kinetic approach to computational geometry, Proc. 24th IEEE Symp. on Foundations of Computer Science, 1983, 100–111.
8. H. Imai, M. Iri, and K. Murota, Voronoi diagram in the Laguerre geometry and its applications, Tech. Rept. RMI 83-02, Dept. of Mathematical Engineering and Instrumentation Physics, University of Tokyo.
9. D. G. Kirkpatrick, Optimal search in planar subdivisions, *SIAM J. Computing* 12 (1983), 28–35.
10. D. Leven and M. Sharir, An efficient and simple motion planning algorithm for a ladder moving in two-dimensional space amidst polygonal barriers, Tech. Rept. 84-014, The Eskenasy Institute of Computer Science, Tel Aviv University, November 1984.
11. T. Lozano-Perez and M. Wesley, An algorithm for planning collision-free paths among polyhedral obstacles, *Comm. ACM* 22 (1979), 560–570.
12. E. E. Moise, *Geometric Topology in Dimension 2 and 3*, Springer-Verlag, New York, 1977.
13. J. Nievergelt and F. P. Preparata, Plane sweeping algorithms for intersecting geometric figures, *Comm. ACM* 25 (1982), 739–747.
14. C. Ó'Dúnlaing, M. Sharir, and C. Yap, Generalized Voronoi diagrams for a ladder: I. Topological analysis, Tech. Rept. 139, Computer Science Dept., Courant Institute, November 1984.
15. C. Ó'Dúnlaing, M. Sharir, and C. Yap, Generalized Voronoi diagrams for a ladder: II. Efficient construction of the diagram, Tech. Rept. 140, Computer Science Dept., Courant Institute, November 1984.
16. C. Ó'Dúnlaing and C. K. Yap, A 'retraction' method for planning the motion of a disc, *J. of Algorithms* 6 (1985) 104–111.
17. T. Ottmann, P. Widmeyer, and D. Wood, A fast algorithm for boolean mask operations, *Inst. f. Angewandte Mathematik und Formale Beschreibungsverfahren, D-7500 Karlsruhe*, Rept. No. 112, 1982.
18. J. T. Schwartz and M. Sharir, On the "Piano Movers" problem: I. The case of a two dimensional rigid polygonal body moving amidst polygonal barriers, *Comm. Pure and Appl. Math.* 36 (1983), 345–398.
19. M. Sharir, Intersection and closest-pair problems for a set of planar discs, *SIAM J. Computing* 14 (1985), 448–468.
20. S. Sifrony and M. Sharir, A new efficient motion planning algorithm for a rod in two-dimensional polygonal space, Tech. Rept. 40/85, The Eskenasy Institute of Computer Sciences, Tel Aviv University, August 1985.
21. G. T. Toussaint, On translating a set of spheres, Tech. Rept. SOCS-84.4, School of Computer Science, McGill University, March 1984.
22. E. Welzl, Constructing the visibility graph for n line segments in $O(n^2)$ time, *Inf. Proc. Letters* 20 (1985), 167–172.
23. F. Aurenhammer, Power diagrams: properties, algorithms, and applications, Tech. Rept. F120, IIG, Tech. Univ. of Graz, Austria, 1983.

Received February 21, 1985, and in revised form November 25, 1985.

Note added in proof. Another proof of Corollary 6.2 is given in [23]. In fact, the results of [23] improve the bound for E^d in Remark (2) above to $O(n^{\lfloor d/2 \rfloor})$.