

On the Complexity of a Single Cell in Certain Arrangements of Surfaces Related to Motion Planning*

Dan Halperin†

Department of Computer Science, Tel-Aviv University,
Tel-Aviv 69978, Israel

Abstract. We obtain near-quadratic upper bounds on the maximum combinatorial complexity of a single cell in certain arrangements of n surfaces in 3-space where the lower bound for this quantity is $\Omega(n^2)$ or slightly larger. We prove a theorem that identifies a collection of topological and combinatorial conditions for a set of surface patches in space, which make the complexity of a single cell in an arrangement induced by these surface patches near-quadratic. We apply this result to arrangements related to motion-planning problems of two types of robot systems with three degrees of freedom and also to a special type of arrangements of triangles in space. The complexity of the entire arrangement in each case that we study can be $\Theta(n^3)$ in the worst case, and our single-cell bounds are of the form $O(n^2\alpha(n))$, $O(n^2 \log n)$, or $O(n^2\alpha(n) \log n)$. The only previously known similar bounds are for the considerably simpler arrangements of planes or of spheres in space, where the bounds are $\Theta(n)$ and $\Theta(n^2)$, respectively. For some of the arrangements that we study we derive near-quadratic-time algorithms to compute a single cell.

1. Introduction

In this paper we study certain arrangements of surfaces in three-dimensional space. Some of the arrangements that we study relate to motion planning of certain moving systems with three degrees of freedom. We also study a special arrangement of triangles in space. We derive several combinatorial and algorithmic results for these arrangements and in particular we obtain near-quadratic upper bounds on the maximum combinatorial complexity of any single cell in these arrangements.

* A preliminary version of this paper has appeared in *Proc. 7th ACM Symposium on Computational Geometry*, North Conway, NH, 1991, pp. 314–323.

† Current address: Robotics Laboratory, Department of Computer Science, Stanford University, Stanford, CA 94305, USA.

1.1. Background

The connection between motion planning and arrangements of curves or surfaces has been noted in many recent studies; see, e.g., [Ca], [AS90], and [SS]. For a detailed discussion on this issue see [GS]. Here, we briefly exemplify this connection. Consider, for example, a motion-planning problem for a system with two degrees of freedom moving among obstacles. The problem can be transformed into a two-dimensional configuration space where every point represents a possible placement of the system. In the configuration space, contacts with the obstacles are represented by “constraint curves” and the moving “robot” is shrunk to a moving point. The point robot cannot cross these constraint curves transversally and so, given an initial free placement z of the robot, it is sufficient to compute the single face of the arrangement of constraint curves, which contains z —this is the set of all the free placements of the robot reachable from z via a collision-free motion. For systems with three degrees of freedom, like the ones that we study here, the configuration space is three dimensional and the obstacles are represented by constraint surfaces. Therefore, we are interested in the complexity and computation of a single cell in such three-dimensional arrangements of constraint surfaces.

It was shown by Guibas *et al.* [GSS] (see also [EGS] and [PSS]) that the combinatorial complexity of a single face in an arrangement of n low-degree algebraic curves in the plane is only $O(\lambda_s(n))$ for some constant parameter s that depends on the degree of the curves; here $\lambda_s(n)$ is the nearly linear maximum length of (n, s) Davenport–Schinzel sequences. Guibas *et al.* [GSS] also provide a near-linear-time algorithm to compute a single face (see [GSS] for more details), which has recently been slightly improved in [CEG⁺]. Thus the general motion-planning problem with two degrees of freedom can be considered settled to a satisfactory extent.

This is far from being the case for motion-planning problems with three degrees of freedom and their induced arrangements of constraint surfaces in space, as there are no analogous sharp bounds, in general, on the maximum possible complexity of a single cell in arrangements of low-degree algebraic surfaces in three-dimensional space. Such arrangements may have $\Theta(n^3)$ total combinatorial complexity in the worst case. There are, however, arrangements for which subcubic bounds on the complexity of a single cell are known. A tight bound $\Theta(n)$ is obvious for arrangements of n planes in space. For arrangements of n spheres in space a tight bound $\Theta(n^2)$ was obtained, employing two different methods, in [Au] and in [KLPS]. In all other arrangements where a subcubic upper bound is known, the gap between the lower and upper bounds is wider. Aronov and Sharir [AS90] (see also [AA]) have obtained a bound $O(n^{7/3})$ on the complexity of a single cell (actually on the complexity of all nonconvex cells) in arrangements of n triangles in space. The known lower bound on the complexity of a single cell in an arrangement of triangles is $\Omega(n^2\alpha(n))$. In [HS] we have previously obtained similar bounds ($O(n^{5/2})$, $O(n^{7/3})$) on the complexity of all “interesting” cells (see Definition 4.1 below) for several types of arrangements of surfaces related to the motion planning of certain systems with three degrees of freedom. The known lower bound for the complexity of a single cell in such arrangements is $\Omega(n^2)$.

Note that, throughout our discussion, we are only interested in the free portions of the configuration space, that is, only in those cells of the arrangement of constraint surfaces that represent placements of the moving system where it does not intersect any obstacle.

It has been a prevalent conjecture in recent years that, in analogy with the two-dimensional case, the actual complexity of a single cell in three-dimensional arrangements of surfaces is near quadratic. The only support for this conjecture so far has come from the result for spheres mentioned above, and from a bound $O(n^2\alpha(n))$ on the complexity of the lower envelope of triangles in space [PS]. In this paper we further substantiate this conjecture by proving near-quadratic upper bounds on the complexity of a single cell in the arrangements that we study.¹

1.2. Summary of Results

We identify a set of topological and combinatorial conditions for a collection of m “blue” surface patches and n “red” surface patches. We show that, when these conditions are satisfied, the complexity of any single cell in the arrangement induced by these surface patches is $O(mn \log n + n^2)$. We show that several types of arrangements, some of them related to motion-planning problems, fulfill these conditions. In a few restricted cases we also obtain near-quadratic-time algorithms to compute a single cell.

The bound that we obtain on the complexity of a single cell in the special arrangements of triangles that we study shows that our technique, unlike all the previous approaches to this problem, distinguishes a single cell from all the “interesting” (nonconvex, for triangles) cells. Our bound in this case is $O(n^2 \log n)$ whereas the complexity of all the nonconvex cells in this arrangement can be $\Omega(n^{7/3})$.

From the motion-planning point of view, we study two moving systems. One is the so-called telescopic arm, first studied by Aronov and Ó’Dúnlaing [AO], which has three degrees of freedom and is moving among obstacles in the plane. We continue the study of the arrangements related to this arm initiated in [HS] and obtain an upper bound $O(n^2 \log n)$ on the complexity of a single cell in these arrangements. We devise a deterministic algorithm with running time $O(n^2 \log^2 n)$ to compute a single cell, which is also a “find-path” algorithm for the original motion-planning problem. In this paper we restrict our discussion to the case of point obstacles, which already poses a considerable challenge. We remark on the extension of our result to the case of polygonal obstacles below, at the end of Section 2.

The second system that we study is a robot arm with three rotating links and a fixed base point, moving among point obstacles in the plane. We remark that this three-link arm is a prevailing kinematic substructure of existing robot

¹Slightly after the results reported in this paper had been established, Aronov and Sharir [AS92] obtained a near-quadratic bound on the complexity of a single cell in arrangements of triangles. See Section 6 below.

manipulators and is therefore a natural problem to study in the framework of algorithmic motion planning. The three-link arm has been recently studied by Cox and Yap [CY] in a variant of the motion-planning problem where the obstacles are not known in advance.

In both systems the entire free configuration space can have $\Omega(n^3)$ complexity. Our results are the first nontrivial near-quadratic bounds on the complexity of a single cell in such a three-dimensional configuration space, where the entire free space can be cubic in size. (This is in contrast with “favorable” motion-planning problems, where the entire free configuration space can be shown to have near-quadratic complexity, as in [LS] and [KS].)

In Section 2 we demonstrate our approach on the arrangement related to the motion planning for a telescopic arm. In Section 3 we state a theorem identifying the conditions under which the new bounds can be obtained more generally and then we apply this theorem to a few more types of arrangements. Algorithms for two types of arrangements are presented in Section 4. In Section 5 we obtain a sharper bound on the complexity of a single cell in the arrangement related to the motion planning for a telescopic arm. Finally, some concluding remarks and open problems are given in Section 6.

2. Exposition of the Method—The Telescopic-Arm Arrangement

We present our technique by an example. For simplicity of exposition, we start with an arrangement induced by a telescopic arm moving among point obstacles in the plane (TA-arrangement, for short). We obtain an upper bound on the maximum combinatorial complexity of any single cell in that arrangement. At the end of this section we remark on the extension of this result to the arrangement induced by the motion planning of this arm moving among *polygonal* obstacles in the plane.

The telescopic arm (TA, for short) consists of two links, \overline{op} and \overline{pq} ; see Fig. 1. o is an anchor point. The first link \overline{op} is a telescopic link which can rotate around o , and extend or shrink along its length. The second link \overline{pq} has a fixed length d , and can rotate around p . This system was studied by Aronov and Ó’Dúnlaing [AO] who showed that the configuration space of this arm moving among polygonal obstacles has $\Theta(n^3)$ connected components in the worst case, and obtained an $O(n^3 \log n)$ -time and $O(n^3)$ -space algorithm to compute it.

Let θ denote the angle between \overline{pq} and the positive x -axis. We assume a fixed reference Cartesian coordinate system such that the anchor point o of the arm




Fig. 1. The telescopic arm.




Fig. 2. The L_i 's of a TA-arrangement.

coincides with the origin. In order to construct a cross section of the configuration space at a fixed θ we choose p as a reference point on the telescopic arm. Each position of the arm can be specified as $Z = (X, \theta)$ where X is the position of p . For the purpose of our analysis, we always present X in a rotated coordinate frame within which θ becomes an upward vertical direction.² Let $O = \{o_i: i = 1, 2, \dots, n\}$ be a set of point obstacles in the plane. Each o_i induces a contact surface σ_i , whose cross section at any fixed θ , $L_i(\theta)$, consists of two elements (see Fig. 2):

- (i) A vertical line segment $\overline{p_i q_i}$ of length d , which we denote by V_i , emanating from the point obstacle o_i downward.
- (ii) The half-line lying on the line defined by the origin and by o_i , which starts at o_i and does not contain the origin, which we denote by R_i .

At any fixed θ , $L_i(\theta)$ represents the forbidden placement of the reference point p due to the obstacle o_i . This can be viewed as a natural generalization of the Minkowski difference to the case of hinged robots, namely, $L_i(\theta)$ is the generalized Minkowski difference of the point obstacle o_i and the robot arm at a fixed θ .

As θ varies, the R_i 's, the rays of the L_i 's, rotate around the origin while at the same time the vertical bars, the V_i 's, remain vertical (as if they were hanging loosely from the o_i 's under gravity). Each L_i traces a surface σ_i in $R^2 \times S^1$, and we denote the arrangement of these n surfaces by \mathcal{A} . We partition each σ_i into two subsurfaces: the *red* surface S_i traced by V_i and the *blue* surface T_i traced by R_i . The cross section of \mathcal{A} at a fixed θ is denoted by \mathcal{A}_θ . Without loss of generality, we restrict our discussion to the quadrant $\{(x, y, \theta) | x \geq 0, \theta \in [0, \pi]\}$.

How does the two-dimensional cross section \mathcal{A}_θ of \mathcal{A} change as θ varies? As the coordinate system rotates, \mathcal{A}_θ changes continuously, but its combinatorial structure remains unchanged, unless one of the following two types of critical events occurs at θ :

- I. A vertex of one V_i meets the vertical bar or the ray of another L_j .
- II. Two vertical bars of two L_i 's overlap.

² Instead of using the parametrization (x, y, θ) we use a transformed representation $(x \sin \theta - y \cos \theta, x \cos \theta + y \sin \theta, \theta)$. For any θ , the first two transformed parameters represent rotating the plane by $\pi/2 - \theta$ that creates the desired effect of making the second link point in an upward vertical direction.

It can easily be shown that the overall number of events of either type in the θ -range $[0, \pi]$ is $O(n^2)$. For more details see [HS].

In [HS] we have obtained an upper bound $O(n^{7/3})$ on the complexity of all the three-dimensional cells in \mathcal{A} that contain a portion of the one-dimensional boundary of a surface in their closure, which clearly serves as an upper bound on the maximum combinatorial complexity of any single cell in the arrangement (cells that do not meet such a boundary have much lower complexity). Our goal here is to obtain an improved bound on the complexity of a single cell.

We assume that the point obstacles are in general position. In particular this means that no three point obstacles are collinear and no two points are coradial (i.e., no two points lie on the same ray from the origin). Under this assumption, we can make the following simple observations:

Observation 2.1. *The blue surfaces in the arrangement \mathcal{A} are pairwise disjoint.*

Observation 2.2. *Each pair of red surfaces may intersect in at most one connected curve and no triple of red surfaces meet at a point.*

Observation 2.1 follows from the fact that the rays R_i do not change their relative position as θ changes, so it suffices to show that the rays do not intersect in one specific cross section \mathcal{A}_θ . However, this is true because we assumed no two point obstacles lie on the same ray from the origin.

For two red surfaces to intersect, their generating vertical bars must overlap at some θ , and such an overlap implies that they become collinear. Since θ is restricted to the range $[0, \pi]$ this can happen at most once (assuming that no critical event occurs at $\theta = 0$) and the resulting intersection is clearly connected. For three red surfaces to meet at a point, the three corresponding vertical bars should become collinear, meaning that the three corresponding point obstacles must be collinear, contradicting our assumption. This implies Observation 2.2.

Let \mathcal{A}_i be the two-dimensional arrangement on the red surface S_i traced by V_i , consisting of red and blue curves corresponding to the intersection of S_i with the other red and blue surfaces. Let r_i (b_i) be the number of red (blue) curves in \mathcal{A}_i . We claim that

Lemma 2.3. *The maximum complexity of k distinct faces in \mathcal{A}_i is $O(r_i + b_i + k)$.*

Proof. Since no three red surfaces meet at a point, the r_i red curves in \mathcal{A}_i are pairwise disjoint. Also, since the blue surfaces are pairwise disjoint, clearly their intersections with S_i are also disjoint. (See Fig. 3; for a fixed θ , the coordinate \bar{y} measures the distance of the intersection of V_i and another bar or ray from the bottommost point of V_i). Therefore, the complexity of k faces in the red (blue) subarrangement of \mathcal{A}_i is at most $O(r_i)$ ($O(b_i)$). Overlaying the two subarrangements and using the combination lemma for planar arrangements of curves (see Lemma 4.1 [GSS]), we obtain the stated bound. \square




Fig. 3. The arrangement \mathcal{A}_i .

The external zone of \mathcal{A}_i , that is, all the faces bordering on the relative boundary of S_i , can be regarded as the single outer face of the arrangement; thus we have

Corollary 2.4. *The complexity of the external zone of \mathcal{A}_i is $O(r_i + b_i)$.*

By Observation 2.2 there are at most $n - 1$ (connected) red curves in \mathcal{A}_i . It is also not difficult to see that a blue surface and a red surface meet along at most one connected simple curve, therefore there are at most n blue curves in \mathcal{A}_i . Plugging these values into the bound of Lemma 2.3, it follows that the complexity of k faces in \mathcal{A}_i is $O(n + k)$. Moreover, since this analysis is true for the two-dimensional arrangement on any red surface we have

Corollary 2.5. *The complexity of any t distinct faces in all the arrangements \mathcal{A}_i together is $O(n^2 + t)$.*

Corollary 2.5 distills the major distinctive feature of the arrangements for which we obtain an improved bound and enables us to apply the proof technique that we use below, namely, the so-called “combination lemma.” It is not true for three-dimensional arrangements of surfaces in general. Indeed, an arrangement of n triangles in space where the complexity of n^2 distinct faces (of the two-dimensional arrangements formed by intersecting each triangle with all the other triangles) is $\Omega(n^{7/3})$ can be constructed; see, e.g., [AS90].

Next, we bound the number of *exposed vertices* in a TA-arrangement \mathcal{A} . An exposed vertex is an intersection point of the relative boundary of one surface with another surface. An exposed vertex is necessarily an endpoint of the intersection curve of a pair of surfaces. Since, as already noted, any pair of surfaces in \mathcal{A} meets along at most one simple curve, there are at most $O(n^2)$ such curves in \mathcal{A}

and therefore

Observation 2.6. *The total number of exposed vertices in a TA-arrangement \mathcal{A} is $O(n^2)$.*

We are now ready to prove the following:

Theorem 2.1. *The maximum combinatorial complexity of any single cell in the arrangement \mathcal{A} related to the motion planning for a telescopic arm moving among n points in the plane is $O(n^2 \log n)$.*

Proof. Let us assume that the cell C whose complexity we wish to bound is designated by a point π in space contained in its interior. In a three-dimensional arrangement the number of vertices in any cell is evidently a good measure for the asymptotic complexity of that cell. There are two types of vertices in our arrangement: the exposed vertices and the nonexposed vertices that we refer to as *internal* vertices. Observation 2.6 implies that we can ignore the contribution of exposed vertices. As to internal vertices, note that each such vertex is the meeting point of two red surfaces and one blue surface. This means that each internal vertex appears both on a red surface and on a blue surface. Consequently we can restrict our discussion to the vertices on a monochromatic collection of surfaces. We choose the red surfaces which enable us to exploit the special property mentioned in Corollary 2.5, stating that, in order to count the number of *vertices* on the boundary of the cell C it suffices to bound the number of red *faces* on the boundary of C .

To bound the number of red faces in C , we apply a “combination lemma” in the following way. (For other “combination lemmas” see, e.g., [EGS], [GSS], and [AS90].) Assume we have m blue surfaces and n red surfaces.

Let us recolor, for a moment, half of our red surfaces green and the other half orange. We start with an arrangement of all the blue surfaces and all the green surfaces and then add the orange surfaces one by one, and bound the maximum possible total increase W in the number of *green* faces in C after we add the orange surfaces; we do not count the newly added orange faces of C . Then we do the symmetric operation of starting with all the blue and all the orange surfaces, adding the green surfaces and bounding the increase in the number of *orange* faces of C after adding the green surfaces; by symmetry, its maximum possible value will also be W .

Let $K(m, n)$ be the number of red (now green or orange) faces bounding a cell C in an arrangement of m blue and n red surfaces. Clearly,

$$K(m, 1) = O(m)$$

and the combination process yields the following recurrence

$$K(m, n) \leq 2K\left(m, \frac{n}{2}\right) + 2W.$$

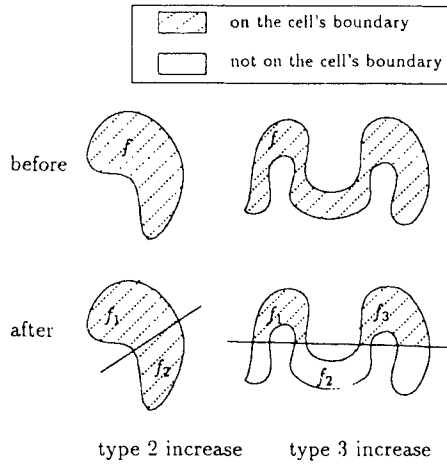


Fig. 4. Various types of increase in the number of green faces.

How large can W be? Suppose we add an orange surface S_i and wish to bound the increase in the number of green faces that belong to ∂C (the boundary of C) due to the addition of S_i . Every green face f that is cut by S_i is represented by an edge (or edges) on the arrangement \mathcal{A}_i of S_i (namely, the connected components of $f \cap S_i$). An increase in the number of green faces can occur in one of the following cases:

1. The edge e of \mathcal{A}_i splitting the green face is an edge of the external zone of \mathcal{A}_i .
2. e is an edge of an internal face of \mathcal{A}_i that cuts a green face f such that both parts of f , in the vicinity of e , still belong to ∂C (as in the left-hand side of Fig. 4).
3. The intersection of the surface S_i with a green face f consists of several edges, cutting f into several subfaces such that some of these subfaces do not belong to ∂C any more but more than one subface still belongs to ∂C (as sketched in the right-hand side of Fig. 4— f_1 and f_3 still belong to ∂C but f_2 no longer belongs to ∂C).

The number of edges causing increase of the first type, over all orange surfaces S_i , is clearly at most $O(mn + n^2)$, by applying Corollary 2.4 to all the arrangements \mathcal{A}_i on the surfaces S_i .

Internal faces that induce an increase of the second type are called in [AS90] *cutting-but-not-splitting*. Aronov and Sharir prove the following topological lemma. They prove it for triangles, but their proof is topological and does not exploit any special property of triangles (beyond those stated below), and so we rephrase the lemma in a more general setting.

Lemma 2.7 (adapted from [AS90]). *Given n surface patches in general position in 3-space, each of which is a simply connected orientable 2-manifold (with boundary), such that each pair of surface patches intersect in at most some fixed*

constant number of connected, simple, open curves, any three surface patches meet in at most one point, and the total number of curves (pairwise intersections of surface patches) is t , then the maximum number of faces that cut-but-do-not-split in an incremental construction of the three-dimensional arrangement of the surfaces is at most $O(t)$.

By our discussion so far, the number of pairwise intersection curves, i.e., the total number of curves in all the arrangements \mathcal{A}_i , is at most $O(mn + n^2)$. Hence this is also a bound on the number of cutting-but-not-splitting faces. (Note that Lemma 2.7 refers to the number of faces that cut-but-do-not-split in the entire arrangement, whereas we only need the bound for a single cell.) By Corollary 2.5 this is also a bound on the number of edges bounding these faces, which is thus an upper bound on the number of increases of type 2.

Remark 2.8. As a matter of fact, in Corollary 2.5 and in Observation 2.6 we have not distinguished between the number of blue surfaces and the number of red surfaces. It can be easily verified that if we denote the number of blue (resp. red) surfaces by m (resp. n) there, then the bound of Corollary 2.5 is $O(mn + n^2 + t)$ and the bound in Observation 2.6 is $O(mn + n^2)$. Consequently, the increase W in the number of red faces of types 1 and 2 is at most $O(mn + n^2)$.

Consider next the third type of increase in more detail. We show:

Lemma 2.9. *The maximum increase of type 3 in the number of green faces is $O(mn)$.*

Proof. Let f be a green face that undergoes an increase of type 3 when the new orange surface S_i is added. Let S_j be the green surface containing f and let γ be the curve of intersection between S_i and S_j . In a type 3 increase we are concerned with a face f that is split by the curve γ into three or more subfaces such that some of them no longer belong to the cell C but at least two subfaces still do. See Fig. 4.

A green surface S_j has a two-dimensional arrangement of green and blue curves on it, formed by intersecting S_j with the other green and blue surfaces. Before inserting the orange surfaces, we preprocess each green surface as follows: We extend a horizontal line (a line parallel to the \bar{y} -axis) from each endpoint of a blue curve until it either reaches the boundary of the green surface or it hits another blue curve. See Fig. 5. Note that some of the extensions already exist, when they are colored green. Note also that, as we add orange surfaces, the horizontal extensions are not split further—the intersection of a green surface and an orange surface is horizontal and we assumed that no three (originally) red surfaces meet at a point. By this preprocessing we have increased the overall number of green faces by at most $O(mn)$, as there are most $O(mn)$ blue curves on green surfaces. Since the blue curves on any green surface are θ -monotone it follows that any green face (after the addition of the horizontal extensions) is θ -monotone, i.e., the intersection of any horizontal line (parallel to the \bar{y} axis) and a green face consists




Fig. 5. "Horizontal" decomposition through the endpoints of blue curves.

of at most one segment. This in turn implies that the addition of an orange surface, whose intersection with the green surface S_j is a horizontal line segment, cannot split a face into more than two subfaces. Hence, beyond the potential increase in the number of green faces due to the preprocessing step, there can be no more increase of the third type. \square

Thus the total increase W of types 1, 2, and 3 is $O(mn + n^2)$.

Finally, we solve the recurrence

$$K(m, 1) = O(m) \quad (1)$$

and

$$K(m, n) \leq 2K\left(m, \frac{n}{2}\right) + O(mn + n^2), \quad (2)$$

to obtain

$$K(m, n) = O(mn \log n + n^2). \quad (3)$$

Originally, we have $m = n$, so the number of red faces bounding the cell C in the whole arrangement is $O(n^2 \log n)$. By Corollary 2.5, this is also a bound on their total complexity, which, by the arguments given above, completes the proof. \square

Remarks. (1) In [Ha] we extend this result to the case of a telescopic arm moving among *polygonal* obstacles having n corners in total, and obtain the same asymptotic bound for this case as well. The proof for the polygonal obstacles requires a series of technical lemmas taking advantage of the special geometry of the TA-arrangement and therefore we have chosen to omit it in this paper and

concentrate on the characteristics of a TA-arrangement that are of more general nature as expressed in the main theorem of the paper, Theorem 3.1.

Briefly, the proof for the case of polygonal obstacles in [Ha] proceeds as follows. We start with the arrangement \mathcal{A} induced by the problem of moving the arm among the corners (vertices) of the polygons, which can be viewed as point obstacles and therefore we can use the bound on the complexity of a single cell which is obtained in Theorem 2.1 above. We then insert into the arrangement \mathcal{A} the extra surfaces induced by the “walls” of the polygonal obstacles (the edges of the polygons), and bound the maximum possible increase in the complexity of a single cell due to this insertion. Exploiting the special structure of these additional surfaces, we show that the overall number of new features showing up on the boundary of a cell due to their insertion is $O(n^2 + f(n))$, where $f(n)$ is the maximum complexity of any cell in \mathcal{A} . Hence, the maximum complexity of a single cell in the final arrangement is $O(n^2 \log n)$ as well.

(2) In Section 5 we improve the bound obtained in Theorem 2.1 for a TA-arrangement. The proof there relies heavily on the specific geometric structure of the TA-arrangement, is much more complex than the proof of Theorem 2.1, and does not seem to generalize easily.

3. Generalization and More Applications

The main result of this paper is stated as Theorem 3.1 in Section 3.1, which abstracts and generalizes the proof of Theorem 2.1. We then present two more applications of this result. In Section 3.2 we analyze the complexity of a single cell in the arrangement related to the motion planning of a three-link planar arm, and in Section 3.3 we consider special arrangements of triangles in space.

3.1. Generalization

Before stating the theorem we define a certain property of a two-dimensional arrangement defined on one surface with respect to a collection of other surfaces. We use the term *surface* here to refer to a simply connected orientable 2-manifold in 3-space with boundary.

Definition 3.1. Let S_i be a surface in 3-space. Let R be a collection of surfaces in 3-space. Let \mathcal{A}_{S_i} be a two-dimensional arrangement of curves on S_i . We say that \mathcal{A}_{S_i} is a *monotone subdivision with respect to R* if the intersection of any face in \mathcal{A}_{S_i} with any surface in R consists of at most one connected component (curve).

This definition, as employed in Theorem 3.1 below, captures the conditions that enable us to bound face increases of the third type. In all the applications in this paper it is almost immediate to show “monotonicity” when necessary. We remark that there is an alternative way to bound face increases of the third type, devised

by Huttenlocher *et al.* [HKS, Proposition 1], when the surfaces are convex polygons in space.

Theorem 3.1. *Let R be a set of n red surfaces and let B be a set of m blue surfaces in space, forming a three-dimensional arrangement \mathcal{A} , such that:*

- (i) *The blue surfaces are pairwise disjoint.*
- (ii) *The intersection between a pair of red surfaces, or between one red and one blue surface, consists of at most some fixed constant number of connected, simple, open curves.*
- (iii) *No three red surfaces meet at a common point.*
- (iv) *For every red surface $S_i \in R$, the arrangement \mathcal{A}_{S_i} of blue curves on S_i (formed by intersecting S_i with all the blue surfaces) can be made into a monotone subdivision with respect to $R \setminus \{S_i\}$ by adding at most $O(m)$ edges to the arrangement \mathcal{A}_{S_i} (see Definition 3.1 above) and any surface in $R \setminus \{S_i\}$ does not meet more than one added edge (of the $O(m)$ edges that refined the arrangement).*

Then the maximum combinatorial complexity of any single cell in \mathcal{A} is at most $O(mn \log n + n^2)$.

Proof. Most of the proof ingredients have already appeared in the proof of Theorem 2.1. Indeed, let $R = \{S_1, S_2, \dots, S_n\}$ be a collection of red surfaces and for each S_i let \mathcal{A}_i be the two-dimensional arrangement of curves formed by intersecting S_i with all the other surfaces in R and B (the collection of blue surfaces). Let r_i be the number of red curves in \mathcal{A}_i , i.e., intersection curves between S_i and other red surfaces, and let b_i be the number of blue curves in \mathcal{A}_i . Then, Observations, Lemmas, and Corollaries 2.3–2.7 of Section 2 still hold (almost) verbatim (see also Remark 2.8). Consequently, the increase W in the number of red faces of types 1 and 2 is at most $O(mn + n^2)$.

To bound the increase of type 3 consider condition (iv) and Definition 3.1 above. Condition (iv) allows us to generalize Lemma 2.9 of Section 2 in the following way. Consider a red surface $S_i \in R$ and the two-dimensional arrangement \mathcal{A}_{S_i} of blue curves formed by intersecting S_i with all the blue surfaces in B . Add $O(m)$ edges to \mathcal{A}_{S_i} to make it monotone with respect to $R \setminus \{S_i\}$; this is possible by condition (iv) of the theorem. Repeat this process for every surface in R . Since no three red surfaces meet at a point, the addition of the intersection curves of S_i with any subset $\bar{R} \subset R$ to the arrangement \mathcal{A}_{S_i} does not violate the monotonicity of the “refined” arrangement with respect to $R \setminus \bar{R}$, because a pair of red curves cannot meet on S_i which is itself a red surface.

The fact that any surface in $R \setminus \{S_i\}$ does not meet more than one added edge assures us that the overall increase in the number of faces in the two-dimensional arrangement on S_i due to the additional $O(m)$ edges is at most $O(m)$. (We allow a surface in $R \setminus \{S_i\}$ to cross one added edge because in the arrangements induced by the motion of robot arms any red surface is connected to a matching blue surface along a curve and an added edge incident to the endpoint of a blue curve on S_i may partially overlap with the intersection of the corresponding red surface.)

Hence, the overall increase in the number of red faces on all the red surfaces, due to the partitioning that made them monotone, is $O(mn)$.

Once we have partitioned all the green surfaces (i.e., half of the red surfaces in one recursive step) in this manner, the intersection of any orange surface with a green face consists of at most one connected component and hence a type 3 increase can no longer occur. Thus, the maximum possible increase of type 3 in the number of green faces in the incremental process is bounded by $O(mn)$.

Consequently, the recurrence relation (1), (2) and its solution (3) are exactly as in the proof of Theorem 2.1. \square

We will show that several naturally defined arrangements of surfaces in space comply with the conditions of the theorem. It is worth noting that the theorem does not impose any constraints on the exact geometry of the surfaces; this is especially useful in the next section where some of the surfaces are obtained by rotating “circular conchoids” in space.

3.2. The Three-Link Arm Arrangement

In this section we consider a standard three-link anchored arm in the plane, which has three rotational degrees of freedom: θ_1 , θ_2 , and θ_3 (see Fig. 6). We number the links according to their order in the kinematic chain beginning from the anchor point. This anchor point is denoted by o ; p is the joint between link 1 and link 2; q connects links 2 and 3; r is the free endpoint of the last link. The length of link i is denoted by d_i . We analyze the arrangement related to the motion planning for this arm moving among point obstacles.

We designate a possible placement of the three-link arm by (X, θ_3) where X is the placement of the joint q in the plane and θ_3 is the angle between link 3 and the positive x -axis. This still falls a little short of a complete specification of the arm configuration (because there can be two distinct placements of the arm with the same (X, θ_3) parametrization) but we overcome this problem below. We fix θ_3 and select the point q connecting links 2 and 3 to be a reference point. For each point obstacle we compute the “generalized Minkowski difference” of the point obstacle and the arm, that is, the locus of all forbidden placements of q due to that point obstacle, for the fixed θ_3 . The Minkowski difference consists of three “constraint” curves, one for each link. We rotate the coordinate axes, using the




Fig. 6. The three-link arm.




Fig. 7. The kinematics of the two-link arm.

same transformation as in Section 2, to make θ_3 always an upward vertical direction; thus the constraint curve induced by the third link is, as before, a vertical segment emanating from the point obstacle downward.

To understand the curves induced by the first two links, we review the kinematics of the two-link planar arm, considering only the case where $d_2 \leq d_1$ (the case $d_2 > d_1$ can be treated similarly). The workspace of this arm (that is, the locus of all reachable placements of the free endpoint q of the second link) is a closed ring of radii $d_1 - d_2$ and $d_1 + d_2$ about the origin o (see Fig. 7(a)). Every point on the boundary of the ring (either on the inner or the outer circle) is attainable by exactly one configuration of the arm in which the links become collinear (Fig. 7(b)). Inside the ring, on the other hand, every point is attainable by two configurations of the arm. When o , p , and q are oriented counterclockwise we call this configuration an *elbow-right* configuration (the solid line configuration in Fig. 7(c)) while an *elbow-left* configuration denotes a configuration where o , p , and q are oriented clockwise (the dashed line configuration in Fig. 7(c)). The degenerate configuration in which o , p , and q are on the same line is arbitrarily considered elbow-right.

Our analysis concentrates on the surfaces induced by elbow-right configurations, so from this point on we ignore the other type of configurations. The analysis of the elbow-left surfaces is completely symmetric. These two sets of surfaces clearly do not interact. Note that for the motion-planning problem, if we are given the initial placement of the arm in say, elbow-right configuration, we may also have to consider the outer cell of the elbow-left arrangement. This, however, does not affect the asymptotic complexity results, because, in any case, we need to consider at most one cell in each arrangement.

Assuming that a point obstacle o_i is at distance $> d_1$ (and $< d_1 + d_2$) from the origin (the anchor point), then the constraint curve due to o_i and link 2 of the arm is the “circular conchoid” traced by q as p rotates about o and pq passes through o_i (Fig. 8). This curve starts at o_i and ends on the *outer* circle of the workspace of




Fig. 8. A conchoidal constraint curve.




Fig. 9. Close-to-origin point obstacles and their induced semicircular constraint curves.

the first two links. The curve lies on one side of the line segment that connects its endpoints.

Consider next, point obstacles whose distance from the anchor point of the arm is $\leq d_1$. Each such point induces a constraint semicircle (traced as the second link rotates around a fixed position of p , in which the first link is fixed and touches the obstacle) that stretches from the inner circle of the workspace ring to the outer circle of that ring, as in Fig. 9. Such a point obstacle may also induce another constraint curve, if it is at least $|d_1 - d_2|$ away from the origin. This curve is also a “circular conchoid” that starts at the point obstacle and ends on the *inner* circle of the workspace.

In summary, under the assumptions we have made so far, the generalized Minkowski difference of a point obstacle o_i and the three-link arm, for a fixed θ_3 , consists of at most three curves: a vertical segment, a circular conchoidal curve, and a circular arc. We denote this combination by $L_i = L_i(\theta_3)$. In addition, for every θ_3 , the cross-section of constraint surfaces is the arrangement formed by the set $\{L_i(\theta_3) | i = 1, \dots, n\}$, where $L_i(\theta_3)$ is the generalized difference of $o_i(\theta_3)$ and the arm, and $o_i(\theta_3)$ is an upward vertical direction. Finally, as θ_3 varies, the two-dimensional arrangement of the conchoidal curves and circular arcs rotates rigidly around the origin whereas the vertical segments remain vertical, and may occasionally swap their horizontal ordering and intersect the other rotating curves.

Back to the two-dimensional arrangement induced by the 2-link arm, there are three families of curves in it: conchoidal curves that end on the outer circle, conchoidal curves that end on the inner circle, and circular arcs. We wish to bound the complexity of a single face in this arrangement, using the following idea by Sifrony [Si]:

Lemma 3.2 [Si]. *Let C be a collection of n Jordan arcs, each pair of which intersect in a constant number of points, and suppose that C can be divided into $k = O(1)$ subsets C_1, \dots, C_k such that each pair of curves within the subset C_i intersect at most s_i times. Then the maximum complexity of a single face in the arrangement induced by C is $O(\lambda_{s+2}(n))$, where $s = \max\{s_1, \dots, s_k\}$.*

In order to obtain a sharp bound on the complexity of a single face in the 2-link-arm arrangement we take an additional precaution and we partition the family of conchoidal curves that end on the outer circle into two subfamilies. For




Fig. 10. A pair of point obstacles inducing conchoidal curves that intersect twice.

a point obstacle o_i that is more than d_1 away from the origin and that induces a conchoidal curve that ends on the outer circle, we mark on the circle of radius d_1 around the origin, the point ξ where the tangent to the circle through o_i touches the circle. We split the conchoidal curve induced by o_i , at the point for which the endpoint p of the generating segment (the second link) touches ξ . This way we have split the conchoidal curve into two: one portion is generated when the angle $\angle opq$ of the arm opq is acute and the other portion is generated when the angle $\angle opq$ is obtuse. We repeat the process for each point obstacle o_j that is farther than d_1 away from the origin. Thus we have split the conchoidal curves that end on the outer circle into subcurves that fall into two families: the “acute” and the “obtuse.” The purpose of the above process is to ensure that a pair of curves within one family intersect at most once. Indeed, a pair of conchoidal curves that end on the outer circle may meet at most twice, but if this is the case then one meeting point is of two acute portions and one meeting point is of two obtuse portions, as is illustrated in Fig. 10.

Next, we employ Lemma 3.2 with the four resulting families of curves and with $s = 1$, to deduce that the maximum complexity of a single face in the 2-link-arm arrangement is $O(n\alpha(n))$ (and that it can be constructed in time $O(n\alpha(n) \log^2 n)$ [GSS], or in expected time $O(n\alpha(n) \log n)$ by a randomized incremental algorithm [CEG⁺]). Consider the two-dimensional cross section of the three-dimensional arrangement, which contains the point π designating the desired cell C and let g be the face containing π of the subarrangement of this cross section which is induced by the first two links; g has $O(n\alpha(n))$ arcs.

Suppose π represents the initial placement of the arm in, say elbow-right configuration. Then we have to consider not only the face g in the two-dimensional arrangement of the elbow-right configuration but also the two external faces h_1 and h_2 of the elbow-left two-dimensional arrangement bordering respectively on the internal and external circles that determine the boundary of the workspace of the 2-link arm. Here also we are interested in the portion of each external face bounded by at most two circular arcs, those arcs that are induced by the same point obstacles that induced the circular arcs in the elbow-right two-dimensional arrangement. It is obvious that the complexity of each of the faces h_1 and h_2 (possibly restricted by one or two circular arcs) is $O(n\alpha(n))$ as well. We restrict our description to handling the portion of the three-dimensional arrangement whose cross section is g . The other relevant parts are handled similarly.

We let g rotate around the origin as θ_3 changes. During this sweep we split each arc into two subarcs at its point of vertical tangency, if it exists (that point depends on θ_3). Let B be the collection of surfaces traced by the arcs of g as θ_3 changes. Now B will play the role of the set of blue surfaces in Theorem 3.1 and the surfaces traced by the vertical bars will be the red surfaces (the collection R) in the theorem. To comply completely with the conditions of the theorem we should make a gap between each pair of blue surfaces created by arcs of g sharing a vertex of g . Note that splitting the arcs of g along points of vertical tangency was necessary to obtain condition (iv) of the theorem. Indeed, the splitting causes the blue intersection curves on any red surface to become θ -monotone—because a red surface is generated by “sweeping” a vertical segment. Therefore, the argument about a type 3 increase is the same as for a telescopic arm (see Lemma 2.9). In this application we have $m = |B| = O(n\alpha(n))$. Thus we can state

Theorem 3.2. *The complexity of a single cell in the three-dimensional arrangement related to the motion-planning problem of a three-link arm moving among n point obstacles in the plane is $O(n^2\alpha(n) \log n)$.*

3.3. Special Arrangement of Triangles in Space

In [AS90] Aronov and Sharir consider an arrangement of n vertical rectangles with two sides parallel to the z -axis, and horizontal convex plates with a total of q vertices. They show that the complexity of a single cell in such an arrangement is $O(n^2\alpha(n) + q)$. We consider a variant of this problem where the horizontal plates are replaced by an arbitrary collection of m pairwise disjoint triangles, and the vertical rectangles are replaced by a collection of n (possibly intersecting) triangles parallel to the z -axis. Otherwise, the triangles are in general position, e.g., no three triangles meet at a segment (in particular, no three vertical triangles intersect), a vertex of one triangle does not lie on another triangle, etc.

Note that in this arrangement \mathcal{A} of n vertical triangles and m other pairwise disjoint triangles, the lower bound on the complexity of a single cell is $\Omega(mn\alpha(n) + n^2)$. Simply construct a collection of $n/2$ segments in the plane whose lower envelope has complexity $\Omega(n\alpha(n))$ [WS], erect a vertical very long triangle from every segment, and then cut this structure by m horizontal triangles. The second term, $\Omega(n^2)$, can be obtained by adding another collection of $n/2$ vertical triangles, each pair of which intersect, outside the preceding structure. It can also be shown that the complexity of all the nonconvex cells of an arrangement of $n/2$ triangles of each type can be $\Omega(n^{7/3})$.

We claim that the two sets of triangles comply with the conditions of Theorem 3.1. Let the pairwise disjoint triangles be the blue surfaces in that theorem and let the triangles parallel to the z -axis be the red ones there. Conditions (i)–(iii) are immediately verified. As for condition (iv), consider a red triangle T (parallel to the z -axis) and the arrangement of blue segments formed on T by intersecting it with all the blue triangles. Extend a line segment parallel to the z -axis from either

endpoint of any blue segment on T until it either reaches an edge of T or hits another blue segment. This results in a refinement of the arrangement on T that is evidently monotone with respect to all the other red triangles and no other red triangle would meet any of the extensions (such a meeting point would imply a degenerate case—an edge of one triangle meets the intersection edge of two other triangles).

Therefore, we have:

Theorem 3.3. *The maximum complexity of a single cell in an arrangement of m pairwise disjoint triangles and n vertical (possibly intersecting) triangles is at most $O(mn \log n + n^2)$.*

If we let $m = n$ in the above Theorem 3.3, then the bound stated in it is $O(n^2 \log n)$. As mentioned earlier, the complexity of all the “interesting” (non-convex) cells in such an arrangement can be $\Omega(n^{7/3})$ in the worst case. This shows that, unlike previous attacks on the single component problem [AS90], [HS], our new approach distinguishes a single cell from all the interesting cells in the arrangement. (The latter technique of [AS92] also distinguishes a single cell from all the interesting cells in the arrangement.)

In the following section we mention a randomized algorithm to compute a single cell in such an arrangement.

4. Algorithms

In this section we present algorithms for computing a single cell for two of the arrangements studied so far in this paper. In Section 4.1 we describe a deterministic algorithm to compute a single cell in the telescopic-arm arrangement described in Section 2. In Section 4.2 we mention a randomized algorithm to compute a single cell in the special arrangement of triangles discussed in Section 3.3. In either case we adapt an existing algorithm to solve our problem of computing a single cell. We therefore omit a detailed description of either algorithm and refer the reader to the original description of these algorithms. Nevertheless, since the adaptation of the algorithm for a TA-arrangement is more involved, we briefly review the main ideas underlying the technique.

4.1. Computing a Single Cell in a TA-Arrangement

We describe a deterministic algorithm for computing a single cell in a TA-arrangement which runs in worst-case $O(n^2 \log^2 n)$ time. In the motion-planning context this algorithm can be viewed as a “find-path” algorithm, i.e., an algorithm that can produce a sequence of elementary motions from the start to the goal placements of the arm, if such a motion is possible. Once the cell is computed, producing the sequence of motions is a fairly standard technique in motion

planning. This algorithm is an adaptation of an existing algorithm for solving related problems. In order to describe the evolution of this algorithm we first provide a couple of necessary definitions.

Definition 4.1. *The interesting cells of an arrangement* are those three-dimensional cells of the arrangement whose boundary contains a portion of the one-dimensional boundary of any surface of the arrangement. For example, as mentioned earlier, the interesting cells in an arrangement of triangles in space are the nonconvex cells in the arrangement. The term is borrowed from [AS90].

Definition 4.2. *A decision algorithm for an arrangement of surfaces in 3-space* is an algorithm that decides whether two given points in 3-space lie in the same connected component (cell) of the given arrangement. If the arrangement consists of “constraint surfaces” of some motion-planning problem (see Section 1), then we call the decision algorithm a *reachability algorithm* as it determines whether there is a continuous collision-free motion between two placements of the “robot” (without necessarily having to produce such a motion).

In [HS] we have presented a decision algorithm for a TA-arrangement (which is based on the reachability algorithm of [HOS]). It was also elaborated there into an algorithm to compute all the interesting cells in a TA-arrangement. The algorithm we present here is a simple modification of the latter. We first give a brief informal description of the previous versions of the algorithm and then describe the necessary modifications so that it will efficiently compute a single cell in a TA-arrangement (for details see [HOS] and [HS]).

The decision algorithm proceeds by sweeping a TA-arrangement \mathcal{A} (in the configuration space as described in Section 2) in the θ direction with a plane parallel to the xy -plane. Throughout the sweep it builds a compact representation of the arrangement as a discrete graph called the *connectivity graph*, CG , where each node represents a certain subcell of \mathcal{A} and two nodes are connected by an edge of CG if there is a direct crossing in \mathcal{A} (not through a surface of \mathcal{A}) between the subcells they represent. The size of the connectivity graph CG is $O(n^2)$. After constructing CG , when given two points, we locate the nodes of CG that represent the subcells containing the points and we look for a path in CG between these nodes. Such a path exists if and only if the given points are contained in the same three-dimensional cell of \mathcal{A} .

However, the information gathered by the decision algorithm is insufficient for computing any cell. To enhance the algorithm into an algorithm to compute all the interesting cells [HS, Section 5.2], we sweep through \mathcal{A} more carefully and whenever we are in an interesting cell we record more information than we did in the original, decision version. The result of this sweep is an enhanced connectivity graph, ECG , containing detailed information about the structure of the interesting cells. If we denote the complexity of all the interesting cells in \mathcal{A} by $g(n)$, then it is shown in [HS] that the algorithm requires $O(n^2 \log^2 n + g(n))$ time and $O(n^2 + g(n))$ space.

We follow the same idea here, namely, we sweep \mathcal{A} as in the decision algorithm but this time we record more information on all the subcells of our single cell C of interest which is designated by some point π in its interior. In order to detect the relevant subcells we have to carry out some preparatory processing. Following is a sketch of the entire algorithm:

1. Run the decision algorithm to produce the succinct connectivity graph CG . Attach to each node of CG information including the smallest θ where the corresponding subcell was detected and the exact geometric data (x, y coordinates) of the event that started the subcell.
2. Exhaustively search CG for the node v representing the subcell which contains the designating point π . Using a breadth-first-search procedure, collect all the nodes of CG that belong to the connected component of v and arrange them in a queue Q by increasing θ order (the same θ value we have attached to any node in step 1).
3. Sweep \mathcal{A} again as in step 1 but for every item in Q enhance CG with more detailed information pertaining to the subcell referred to in the item. (This is done in a way similar to recording the interesting cells in [HS].) This step still leaves some portions of relevant subcells unattended because they exist for a smaller θ than the θ in which the relevance of the subcell is detected. Thus we need one additional step.
4. Reverse the order of the items in Q into a new queue Q^R and run a sweep similar to the previous step, but running backward in θ so that portions of subcells that were not followed by the previous step can be traced now.

The complexity analysis of the above algorithm is similar to that in [HS] and so we omit it. If we denote the maximum complexity of any single cell in \mathcal{A} by $h(n)$ then the algorithm requires $O(n^2 \log^2 n + h(n))$ time and $O(n^2 + h(n))$ space. Thus, by Theorem 2.1 we have

Theorem 4.1. *A single cell in an arrangement \mathcal{A} of surfaces induced by the motion planning for a telescopic arm among n point obstacles in the plane can be computed in $O(n^2 \log^2 n)$ time and $O(n^2 \log n)$ space.*

4.2. An Algorithm for Special Arrangements of Triangles in Space

We apply a randomized algorithm for computing a single cell in an arrangement of triangles in space, devised by Aronov and Sharir in [AS90], to the special arrangement of triangles described in Section 3.3. In [AS90] the authors express the running time of the algorithm as a function of the maximum combinatorial complexity of a single cell in the arrangement. By Theorem 3.3, the bound on this quantity is $O(n^2 \log n)$ in our special case. The application of the algorithm is straightforward, and we only state the final result here (for details, see [AS90] and Section 5.4 of [Ha]).

Theorem 4.2. *Given a fixed $\delta > 0$ and two sets of triangles in space T_1 and T_2 , T_1 consisting of pairwise disjoint triangles and T_2 consisting of triangles all parallel to a common line in space, $|T_1| + |T_2| = n$. A single cell in the arrangement $\mathcal{A}(T_1 \cup T_2)$ can be computed in randomized expected time $O(n^{2+\delta})$, where the constant of proportionality depends on δ .*

5. A Sharper Bound for a TA-Arrangement

In this section we derive an improved bound $O(n^2\alpha(n))$ on the complexity of a single cell in a TA-arrangement. The proof of this bound relies heavily on the specific geometric structure of the TA-arrangement, and is much more complex than the proof of Theorem 2.1. The bound obtained here suggests that the log factor showing up in the results in this paper is an artifact of the proof technique rather than describing the inherent complexity of a single cell in the arrangements that we study.

Theorem 5.1. *The maximum combinatorial complexity of a single cell in a TA-arrangement \mathcal{A} is $O(n^2\alpha(n))$.*

Proof. Without loss of generality we restrict our discussion to the quadrant $\{(x, y, \theta) | x \geq 0, \theta \in [0, \pi]\}$. To bound the complexity of a single cell in \mathcal{A} , it suffices to bound the complexity of the “zone” of the plane $\theta = \theta_0$ in \mathcal{A} . For convenience we assume that $\theta_0 = 0$ (we can always rearrange the coordinate frame so that $\theta_0 = 0$), and denote the plane $\theta = 0$ by $P_{\theta=0}$. The zone of $P_{\theta=0}$ is the collection of cells of zero to three dimensions which are reachable from (the free portion of) $P_{\theta=0}$. We concentrate on one side of $P_{\theta=0}$; the arguments for the other side are symmetric.

As before, we denote the vertical bar of L_i by V_i , its upper vertex by p_i , and its lower vertex by q_i . The length of V_i (that is, the length of the second link of the TA) is denoted by d .

For each L_i we define the following planar arrangement. For every L_j , including $j = i$, we define the function $F_j(\theta)$ to be the x -coordinate of the intersection point between the vertical bar of L_j and the ray of L_i . The collection of the graphs of the functions $F_j(\theta)$ is easily seen to be an arrangement of pseudosegments in the $X\theta$ plane, which we denote by \mathcal{A}_i .

We wish to count the number of edges and vertices in the zone of $P_{\theta=0}$. The only features of \mathcal{A} that do not appear as features of the \mathcal{A}_i 's are those involving the spatial curves traced by the lower endpoints q_i of the vertical bars of the L_i 's. However, these are exposed edges and vertices and we have argued before (Observation 2.6) that the maximum number of all these features in the entire arrangement \mathcal{A} is $O(n^2)$. Therefore, the overall complexity of the collection $\{\mathcal{A}_i: i = 1, 2, \dots, n\}$ dominates the complexity of \mathcal{A} .

Our plan is to modify each arrangement \mathcal{A}_i into an arrangement \mathcal{B}_i so that \mathcal{B}_i will have the following properties:

- P1. Every face of a subcell in \mathcal{A} whose “ceiling” is L_i is represented by an edge of \mathcal{B}_i unless it is accounted for during the modification process.
- P2. Given a face in \mathcal{B}_i that is *not* in the (planar) zone of the line $\theta = 0$ in \mathcal{B}_i , the three-dimensional subcell of \mathcal{A} whose ceiling is this face is inaccessible from the plane $P_{\theta=0}$ in \mathcal{A} . (We define the subcell of \mathcal{A} whose *ceiling* is an inner face f of \mathcal{B}_i to be the collection of points in 3-space that can be connected by a segment parallel to the y -axis and directed in the positive y -direction to a point in f , without crossing any other surface. In other words, the subcell associated with a ceiling f , is the collection of points that can be seen from f when looking in the negative y -direction.) Note that the zone of the line $\theta = 0$ is actually the *outer face* of \mathcal{B}_i (in the infinite $X\theta$ plane) since we restrict \mathcal{B}_i to the θ -interval $[0, \pi]$ and to $x \geq 0$.

The first property assures us that every element we have to count appears in one of the \mathcal{B}_i 's, or is accounted for separately in the modification process. The second property enables us to ignore parts of the planar arrangements as we count.

How do we build \mathcal{B}_i ? Let us go back to \mathcal{A}_i and color each point on every pseudosegment of \mathcal{A}_i as either *blue* or *red* according to the following rule. At a fixed θ , a point on a pseudosegment of \mathcal{A}_i denotes the intersection between the ray of L_i and the vertical bar of some L_j at that θ . If, when we look from the intersection point in the negative y -direction, we see the endpoint of the vertical bar of L_j (i.e., it is not crossed by the ray of another L_k in that direction), then the corresponding point of \mathcal{A}_i is colored blue, otherwise it is colored red. (See Fig. 11 for an illustration.) This rule applies to the vertical bar V_i of L_i as well. The blue points are “weaker” than the red points, in the following sense (to be formalized below): if two faces f_1, f_2 of \mathcal{A}_i have a common blue boundary point, then it is possible, in the three-dimensional arrangement \mathcal{A} , to cross from points whose ceiling is in f_1 to points whose ceiling is in f_2 by passing below the lower endpoint of the vertical bar that created the blue point (see Fig. 11).

As a first step toward producing \mathcal{B}_i , we simply eliminate the blue portions of the pseudosegments of the colored \mathcal{A}_i . Let us denote the resulting arrangement by \mathcal{A}'_i . We refer to this elimination of blue portions as **Step 1**. We later show, in




Fig. 11. The coloring of \mathcal{A}_i .




Fig. 12. Exposed vertical edge.

Lemma 5.5, that at most $O(n^2)$ features (edges and vertices) of all the arrangements \mathcal{A}_i together were eliminated at Step 1.

Now it is our purpose to show that, roughly speaking, most inner faces of \mathcal{A}'_i represent subcells of \mathcal{A} with ceiling L_i which are inaccessible from any other subcell of \mathcal{A} that has a different L_j as ceiling. As a matter of fact, the only way in which one can cross between two subcells with different ceilings is to pass below the lower endpoint of some V_j which is not intersected by any ray (we call such a V_j an *exposed vertical bar*), and such that V_j has the other ceiling right above it. See Fig. 12, where V_j is an exposed vertical bar enabling crossing between two subcells, one having L_i as the ceiling and the other having L_j as the ceiling. The difficulty in handling such a crossing is that the new ceiling (L_i in Fig. 12) has no record of this crossing in its associated arrangement \mathcal{A}'_i .

In what follows we show that such crossings may occur only in faces of θ cross sections whose corresponding face in \mathcal{A}'_i is the outer face of \mathcal{A}'_i or that it may happen in a special set of faces of \mathcal{A}'_i that we are going to connect to the outer face by creating thin passages between some faces of \mathcal{A}'_i . There are three cases of crossings that need to be considered:

Case 1: The "Upper" Case. In this case the (exposed) vertical bar V_j of some L_j is fully contained in a θ cross section of a cell whose ceiling is L_i and the ray from the origin through q_j (the lower vertex of V_j) has *positive* slope. We will show that the corresponding face of \mathcal{A}'_i is actually the outer face.

Consider the face f of the θ cross section of the three-dimensional arrangement where this happens at, say θ_1 . Let R be the ray through the origin and $q_f(\theta_1)$ and suppose, without loss of generality, that the floor of f is a portion of R (see Fig. 13); we explain below why this assumption does not impair our proof. Denote by V the vertical line that contains V_j and let v denote its intersection point with L_i . It is easily verified that there cannot exist any internal vertex p_m of some L_m inside f which is to the left of V . Next draw the path of $v = v(\theta)$ inside \mathcal{A}'_i for θ between θ_1 and the minimum of φ_i and π , where φ_i is the orientation in which L_i becomes vertical. Since L_i has a positive slope and since v is on L_i this path is well defined.




Fig. 14. Case 2, cross section at θ_1 .

rotates). We will show that the graph of $w(\theta)$ drawn in \mathcal{A}'_i never does meet any pseudosegment of \mathcal{A}'_i as θ decreases from θ_1 to $\theta = 0$ and thus, by arguments similar to those in Case 1, the face \hat{f} of \mathcal{A}'_i corresponding to f is the outer face. More precisely, $w(\theta)$ might not be defined all the way down to $\theta = 0$, it might be defined only until $w(\theta)$ reaches the endpoint p_i of L_i . At every θ in its domain, $w(\theta)$ is on L_i , therefore its graph can be drawn in \mathcal{A}'_i . Suppose $w(\theta)$ meets a pseudosegment s of \mathcal{A}'_i at some $\theta_2 < \theta_1$ in its domain of definition. This means that some vertical bar V_k , namely, the one corresponding to s , whose upper endpoint is above L_i , reaches the floor of the subcell at θ_2 vertically below $w(\theta_2)$. However, it is easily checked that the floor below $w(\theta_2)$ cannot intersect the vertical segment $\overline{w(\theta_2)u}$ because then, rotating back to θ_1 , the ray constituting the floor would have to cross $\overline{w(\theta_1)u}$, contrary to the assumption the V_j is exposed (see Fig. 14). Since the vertical “depth” below $w(\theta)$ for $\theta = \theta_1$ is at least d and it increases as θ decreases (here we use the fact that the slope of L_i at θ_1 is negative), we obtain a contradiction, as in Case 1. Thus \hat{f} is indeed the outer face of \mathcal{A}'_i .

Case 3: The “Middle” Case. So far we have dealt with cases where the exposed vertical bar was fully above the x -axis (Case 1) or fully below the x -axis (Case 2). Next we analyze the case in which the exposed vertical bar crosses the x -axis. (See Fig. 15.) The analysis of this case is a bit more intricate as we can no longer ascertain that the relevant face of \mathcal{A}'_i is the outer face; instead we connect these faces of \mathcal{A}'_i to the outer face of \mathcal{A}'_i by making small holes in certain pseudosegments of \mathcal{A}'_i . Let us start the analysis of this case by making some simple observations.




Fig. 15. Case 3.

Observation 5.1. *At any given θ there can be at most one exposed vertical bar V_j that is fully contained in a face f of the θ cross section of \mathcal{A} with ceiling L_i and crosses the x -axis.*

Indeed, let V_j be the leftmost such vertical bar and suppose V_k is another vertical bar fully contained in f . However, V_k is to the right of V_j and the floor of f has positive slope there (it is at least as high as the ray of L_j), thus V_k cannot be fully contained within f . Thus at any θ there is at most one exposed vertical bar under ceiling L_i that crosses the x -axis.

Denote by $x_j(\theta)$ the x -coordinate of the intersection point between V_j and the x -axis, when it is defined.

Observation 5.2. *If at θ_1 the vertical bar V_k is exposed under L_i and crosses the x -axis, and at θ_2 another vertical bar V_l is exposed under the same L_i and crosses the x -axis, and $\theta_2 > \theta_1$, then $x_l(\theta_2) < x_k(\theta_1)$.*

Indeed, at θ_1 the floor under L_i from V_k to the right is at least as high as the ray of L_k which has positive slope and keeps moving leftward and up as θ increases. Hence a later exposed vertical bar that has L_i as the ceiling and crosses the x -axis will have to appear to the left of $x_k(\theta_1)$.

Our plan is to fix L_i and trace all the exposed vertical bars that cross the x -axis under the ray of L_i . At any θ in which there exists such an exposed vertical bar V_j under L_i we define $w(\theta)$ to be the point of L_i vertically above V_j . By Observation 5.1 there is at most one such point for any θ . Again, since $w(\theta)$ is always on the ray of L_i it can also be regarded as tracing a curve C_i in \mathcal{A}'_i . This curve is θ -monotone but not necessarily connected. We next describe how to extend C_i into a connected curve.

Let θ_1 be the smallest θ for which there is an exposed vertical bar V_j that crosses the x -axis under L_i . Our function $w(\theta)$ traces V_j throughout its exposure, that is, $w(\theta)$ is the point on L_i vertically above V_j throughout the corresponding θ interval. When V_j stops being exposed and as long as no other vertical bar becomes exposed, $w(\theta)$ is defined as follows:

- (i) if V_j can still see L_i above it, we define $w(\theta)$ to be the point on L_i directly above V_j ;
- (ii) if at some stage V_j is bypassed from above by some V_k and V_k can see L_i above it we switch to tracing V_k (so now $w(\theta)$ is the point on L_i directly above V_k); if V_k is bypassed in the same way, we switch to tracing the new vertical bar and so on.

V_k might have become an exposed vertical bar, so its tracing now “joins” a new portion of the original curve C_i .

So far the function $w(\theta)$ is defined continuously. A problem arises when the next exposed vertical bar appears somewhere to the left (recall Observation 5.2) of the current vertical bar that we trace (which is not necessarily exposed, but can see L_i). In this case the function $w(\theta)$ makes a “jump” in the x -coordinate. In \mathcal{A}'_i we

Table 1. Summary of notation for the arrangements related with ceiling L_i .

Description	Symbol	Number of pseudosegments
Original arrangement	\mathcal{A}_i	$O(n)$
After Step 1	\mathcal{A}'_i	n'_i
After Step 2	\mathcal{A}''_i	n''_i
After Step 3	\mathcal{B}_i	m_i

simply connect the two endpoints of the graph of $w(\theta)$ by a segment parallel to the x -axis. In terms of $w(\theta)$ we can view this connecting segment as if we let the tracing point on L_i move along L_i in an infinite speed from where it is currently to the point above the newly exposed vertical bar.


The tracing ends when either the ray of L_i becomes vertical or when θ reaches π . Correspondingly, the curve C_i ends in the outer face of \mathcal{A}'_i . Note that by definition (and recalling Observation 5.2), $w(\theta)$ is x -monotone.

This construction is repeated for the ray of each L_i .

To be able to handle Case 3 successfully, we have to alter each arrangement \mathcal{A}'_i still further, in the following two additional steps (refer to Table 1 for a summary of notation):

Step 2. Whenever a vertical bar V_k , whose upper endpoint is above L_i and whose lower endpoint is below L_i and does not see L_i (that is, V_k crosses the floor of a face whose ceiling is L_i), changes the floor that it crosses (in which case V_k partially overlaps the vertical bar V_i corresponding to that floor, as in Fig. 16), make a small gap in the corresponding pseudosegment of \mathcal{A}'_i . This ensures that pseudosegments of \mathcal{A}'_i that involve different floors will be treated as separate entities. Denote the arrangement \mathcal{A}_i after the changes of Steps 1 and 2 by \mathcal{A}''_i , and let n''_i denote the number of its pseudosegments.

Step 3. For every intersection of the curve C_i with a pseudosegment of \mathcal{A}'_i make a small gap in the pseudosegment such that, in the resulting arrangement, C_i does

**Fig. 16.** Change of floors.

not cross any pseudosegment. In the new arrangement, C_i is fully contained in the outer face.

We denote the final arrangement by \mathcal{B}_i , and let m_i denote the number of its pseudosegments.

By our discussion so far, an internal face of \mathcal{B}_i represents a three-dimensional cell of \mathcal{A} whose ceiling is L_i and which is inaccessible from other three-dimensional cells that have another L_j as a ceiling. Hence, to get an upper bound on the complexity of a single cell, it suffices to bound the overall complexity of the outer faces in all arrangements \mathcal{B}_i . Of course, now we have to show that the price we have paid for this simplification is not too high. It will actually suffice to show that the overall number, $\sum_{i=1}^n m_i$, of pseudosegments that we have created, is still not too large. Indeed, we proceed to show this.

We define a *blue arc* to be a maximal blue connected portion of a pseudosegment in an arrangement \mathcal{A}_i . When producing \mathcal{A}'_i from \mathcal{A}_i (Step 1), we eliminate the blue arcs of \mathcal{A}_i . Each elimination of a blue arc may possibly increase by one the number of pseudosegments in \mathcal{A}'_i over their number in \mathcal{A}_i . The following lemma guarantees that the overall increase in pseudosegments is under control.

Lemma 5.3. *There are at most $O(n^2)$ blue arcs in all the arrangements \mathcal{A}_i together.*

Proof. Either endpoint of a blue arc in \mathcal{A}_i is related to a critical event of type I or type II (see Section 2 and [HS]) as follows: A blue arc in \mathcal{A}_i may start at the beginning of a pseudosegment, which is a type I event. It may also start when L_i witnesses a type II event in which the lower vertex of the vertical bar generating the blue arc also participates. Similarly, it may start when the lower vertex of the vertical bar generating the blue arc meets a ray of another L_j . Likewise, the arc must end in a similar event of either type. There is no other way in which a blue arc can start or end. Each critical event may influence the beginning or end of blue arcs in only a constant number of arrangements \mathcal{A}_i (at most two blue arcs in each). Consequently there are at most $O(n^2)$ blue arcs in all the arrangements \mathcal{A}_i together. \square

Let n'_i denote the number of pseudosegments in \mathcal{A}'_i .

Corollary 5.4. $\sum_{i=1}^n n'_i = O(n^2)$.

Proof. Originally, there are $O(n)$ pseudosegments in each \mathcal{A}_i , because the beginning and end of each pseudosegment is associated with a unique critical event in which L_i participates, and L_i participates in at most $O(n)$ critical events. By Lemma 5.3 the total increase in the number of pseudosegments when transforming the arrangements \mathcal{A}_i into the arrangements \mathcal{A}'_i is $O(n^2)$, and the corollary follows. \square

Moreover, it is easy to show, using similar arguments, that the following holds:

Lemma 5.5. *The maximum number of vertices of all the arrangements \mathcal{A}_i incident to all the blue arcs is $O(n^2)$.*

It is easily verified that Step 2 has a similar effect on the number of pseudosegments, that is, we create at most $O(n^2)$ additional gaps in the existing pseudosegments and so we increase their overall number by $O(n^2)$, making the resulting number still $O(n^2)$.

Next we show that Step 3 also does not cause too much damage. We show this by proving that the curve C_i does not meet a pseudosegment of \mathcal{A}_i'' more than once. The curve C_i , which contains the graph of the function $w(\theta)$, consists of two types of arcs, one type is where $w(\theta)$ traces vertical bars under L_i , and the other type is the “jumps,” the x -parallel portions of C_i . Note that the portions of the curve C_i of the first type do not cross any pseudosegments of \mathcal{A}_i'' because if a vertical bar that generates a red pseudosegment e of \mathcal{A}_i meets $w(\theta)$, then it necessarily changes the floor below L_i (because e is red and in these portions $w(\theta)$ is always above a vertical bar of another L_j) and in Step 2 we split those pseudosegments at exactly such points, so there are small holes through which C_i passes without meeting a pseudosegment of \mathcal{A}_i'' .

As for portions of C_i of the second type, we will show that two different “jumps” cannot share the same floor below them. If we show this, then Step 2 implies that no pseudosegment of \mathcal{A}_i'' can be crossed by both these C_i portions. (It is evident that within a single jump, C_i cannot cross the same pseudosegment twice, because the pseudosegments are θ -monotone.) By *floors* we mean all the rays (of some L_f 's) that lie immediately below L_i when the tracing-point moves along L_i during a jump at a given θ to a point above the next exposed vertical bar V_k . Suppose two jumps occur at θ_1 and θ_2 for $\theta_1 < \theta_2$. Denote by J_1 (resp. J_2) the portion of L_i passed by the tracing-point at the jump of θ_1 (resp. θ_2). Clearly, the floors below J_1 at θ_1 have positive slope (because we are moving left toward a vertical bar whose matching ray has positive slope). We claim that rays that have positive slope at θ_1 cannot appear as floors below J_2 at θ_2 . Note that at θ_1 all the rays that have positive slope must be to the right of the newly exposed vertical bar V_k . If they have a slope smaller than that of L_k at θ_1 , then they cannot show up at all at θ_2 . If the ray of a certain L_m has a positive slope at θ_1 which is larger than that of L_k , then in order to appear as a floor at a later jump, the corresponding vertical bar V_m will have to bypass the traced vertical bar at some stage but when it does so we will switch to tracing it (V_m) and its floor will never be seen below the tracing point. To clarify this issue further, we note that if L_m contributes to the floor under J_1 , then the portion of L_m that constitutes part of this floor is necessarily a contiguous portion of L_m including the terminus of the ray of L_m , i.e., the upper vertex p_m of V_m . Moreover, in order for L_m to contribute to the floor under J_2 , p_m must see L_i for every θ between θ_1 and θ_2 , because once a floor vanishes it cannot reappear in the θ range of our concern. Therefore, there is some θ' , $\theta_1 \leq \theta' \leq \theta_2$, for which $w(\theta')$ is right above p_m and the function w will stop tracing p_m only for an exposed vertical bar that crosses the x -axis and is to the left of x_m . However,

recalling that $w(\theta)$ is x -monotone, this is a contradiction because at θ_2 , before executing the “jump,” the function w traces a vertical bar to the right of V_m .

Consequently, we are allowed to cut through all the pseudosegments of \mathcal{A}'_i that C_i traverses, over all arrangements \mathcal{A}'_i , and still increase the overall number of pseudosegments by only $O(n^2)$. Hence, our discussion so far implies

Lemma 5.6. $\sum_{i=1}^n m_i = O(n^2)$.

Property P2 of \mathcal{B}_i implies that the complexity of the outer face of \mathcal{B}_i is an upper bound on the complexity of the subcells of \mathcal{A} whose ceiling is L_i and are accessible from $P_{\theta=0}$. From [GSS] we know that the complexity of a single face in an arrangement of n pseudosegments is $O(n\alpha(n))$. Therefore the complexity of the outer face of \mathcal{B}_i is $O(m_i\alpha(m_i))$. Summing over all the arrangements \mathcal{B}_i and using Lemma 5.6, we get that the total combinatorial complexity of all the subcells of \mathcal{A} in the zone of the plane $P_{\theta=0}$ is $O(n^2\alpha(n))$. This completes the proof of Theorem 5.1. \square

6. Conclusion

In this paper we have obtained near-quadratic upper bounds on the complexity of a single cell in certain types of arrangements of surface patches in space, most of which arise in the context of motion planning for certain systems with three degrees of freedom, namely, the telescopic arm and the 3-link arm. These were the first results of this kind for motion-planning problems with three degrees of freedom where the entire free configuration space could be cubic in size.

In Section 5 we have derived an improved bound $O(n^2\alpha(n))$ on the complexity of a single cell in a TA-arrangement. The proof there relies heavily on the specific geometric structure of the TA-arrangement, is much more complex than the proof of Theorem 3.1, and does not seem to generalize easily. Still, it implies that the log factor showing up in the main result of this paper (Theorem 3.1) is an artifact of the proof technique rather than describing the inherent complexity of a single cell in the arrangements that we study. We continue to support the conjecture that the complexity of a single cell in arrangements of more general “well-behaving” surfaces is $O(n^2\beta(n))$, where $\beta(n)$ is some function of $\alpha(n)$, determined by the intersection properties of the given surfaces.

The paper raises several open problems. The main open problem is to extend the result to more general arrangements of surfaces. A major step in this direction has been recently made by Aronov and Sharir [AS92] who obtain a bound $O(n^2 \log n)$ on the complexity of a single cell in an arrangement of triangles in space. Another candidate system for obtaining a near-quadratic bound has been the one generated by the motion of an L-shaped object in the plane (see [HOS]). Indeed in [Ha] we derive a bound $O(n^2 \log^2 n)$ on the maximum complexity of a single cell in the arrangement corresponding to this system.

Another open problem is to develop general techniques for constructing a single

cell in arrangements that conform to the conditions of Theorem 3.1, which run in close to quadratic time.

Given a collection G of n surfaces, the zone of an additional surface s in the arrangement $\mathcal{A}(G)$ is the collection of cells that are incident to s in $\mathcal{A}(G \cup \{s\})$. In many cases the problem of bounding the zone complexity of an object in an arrangement \mathcal{A} of the same type of objects is closely related to that of bounding the complexity of a single cell in \mathcal{A} (see, e.g., [EGP⁺]) and we believe that the results presented in this paper can be extended to give similar bounds for the zone complexity of “well-behaving” surfaces in the arrangements studied here.

One of the robot arms that we have studied, the three-link arm, is a common kinematic substructure in existing robot manipulators. It is an interesting problem to exploit the analysis in this paper to derive bounds on the complexity of a single cell in the configuration space of spatial robot manipulators that contain the 3-link arm as a substructure, like the four degrees-of-freedom SCARA type robot or the six degrees-of-freedom “Elbow” manipulator (see [Pa]).

Acknowledgments

The author is grateful to Boris Aronov and Micha Sharir for many helpful discussions on the contents of this paper.

References

- [AA] P. K. Agarwal and B. Aronov, Counting facets and incidences, *Discrete and Computational Geometry* 7 (1992), 359–369.
- [AO] B. Aronov and C. Ó’Dúlaing, Analysis of the Motion-Planning Problem for a Simple Two-Link Planar Arm, Technical Report 314, Courant Institute of Mathematical Sciences, 1987.
- [AS90] B. Aronov and M. Sharir, Triangles in space, or building (and analyzing) castles in the air, *Combinatorica* 10 (1990), 137–173.
- [AS92] B. Aronov and M. Sharir, Castles in the air revisited, *Proc. 8th Symposium on Computational Geometry*, 1992, pp. 146–156.
- [Au] F. Aurenhammer, Power diagrams: Properties, algorithms and applications, *SIAM Journal on Computing* 16 (1987), 78–96.
- [Ca] J. Canny, The Complexity of Robot Motion Planning, Ph.D. Dissertation, Computer Science Department, M.I.T., May 1987.
- [CEG⁺] B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Snoeyink, Computing a single face in an arrangement of line segments, *Proc. 2nd SIAM–ACM Symposium on Discrete Algorithms*, 1991, pp. 441–448.
- [CY] J. Cox and C. Yap, On-line motion planning: moving a planar arm by probing an unknown environment, Manuscript, 1989.
- [EGP⁺] H. Edelsbrunner, L. J. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir, Arrangements of curves in the plane: combinatorics, topology and algorithms, *Proc. 15th International Colloquium on Automata, Languages and Programming*, 1988, pp. 214–229.
- [EGS] H. Edelsbrunner, L. J. Guibas, and M. Sharir, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete and Computational Geometry* 5 (1990), 161–196.

- [GS] L. J. Guibas and M. Sharir, Combinatorics and algorithms of arrangements, in *New Trends in Discrete and Computational Geometry* (J. Pach, ed.), Springer-Verlag, New York, 1993, pp. 9–36.
- [GSS] L. J. Guibas, M. Sharir, and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete and Computational Geometry* 4 (1989), 491–521.
- [Ha] D. Halperin, Algorithmic Motion Planning via Arrangements of Curves and of Surfaces, Ph.D. Thesis, Tel-Aviv University, July 1992.
- [HOS] D. Halperin, M. H. Overmars, and M. Sharir, Efficient motion planning for an L-shaped object, *SIAM Journal on Computing* 21 (1992), 1–23.
- [HS] D. Halperin and M. Sharir, Improved combinatorial bounds and efficient techniques for certain motion planning problems with three degrees of freedom, *Computational Geometry: Theory and Applications* 1 (1992), 269–303.
- [HKS] D. P. Huttenlocher, K. Kedem, and M. Sharir, The upper envelope of Voronoi surfaces and its applications, *Proc. 7th Symposium on Computational Geometry*, 1990, pp. 194–203.
- [KLPS] K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete and Computational Geometry* 1 (1986), 59–71.
- [KS] K. Kedem and M. Sharir, An efficient motion planning algorithm for a convex polygonal object in two-dimensional polygonal space, *Discrete and Computational Geometry* 5 (1990), 43–75.
- [LS] D. Leven and M. Sharir, An efficient and simple motion planning algorithm for a ladder moving in two-dimensional space amidst polygonal barriers, *Journal of Algorithms* 8 (1987), 192–215.
- [PS] J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: combinatorial analysis, *Discrete and Computational Geometry* 4 (1989), 291–309.
- [Pa] R. P. Paul, *Robot Manipulators: Mathematics, Programming and Control*, The MIT Press, Cambridge, MA, 1981.
- [PSS] R. Pollack, M. Sharir, and S. Sifrony, Separating two simple polygons by a sequence of translations, *Discrete and Computational Geometry* 3 (1988), 123–136.
- [Si] S. Sifrony, Efficient Algorithms for Motion Planning Problems in Robotics, Ph.D. Thesis, Tel-Aviv University, January 1989.
- [SS] S. Sifrony and M. Sharir, An efficient motion planning algorithm for a rod moving in two-dimensional polygonal space, *Algorithmica* 2 (1987), 367–402.
- [WS] A. Wiernik and M. Sharir, Planar realizations of nonlinear Davenport–Schinzel sequences by segments, *Discrete and Computational Geometry* 3 (1988), 15–47.

Received February 5, 1992, and in revised form August 11, 1992.