

An Efficient Motion-Planning Algorithm for a Convex Polygonal Object in Two-Dimensional Polygonal Space*

K. Kedem¹ and M. Sharir^{1,2}

¹ Computer Science Department, School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

² Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

Abstract. We present an efficient algorithm for planning the motion of a convex polygonal body B in two-dimensional space bounded by a collection of polygonal obstacles. Our algorithm extends and combines the techniques of Leven and Sharir and of Sifrony and Sharir used for the case in which B is a line segment (a “ladder”). It also makes use of the results of Kedem and Sharir on the planning of translational motion of B amidst polygonal obstacles, and of a recent result of Leven and Sharir on the number of free critical contacts of B with such polygonal obstacles. The algorithm runs in time $O(kn\lambda_\delta(kn) \log kn)$, where k is the number of sides of B , n is the number of obstacle edges, and $\lambda_\delta(q)$ is an almost linear function of q yielding the maximal number of connected portions of q continuous functions which compose the graph of their lower envelope, where it is assumed that each pair of these functions intersect in at most s points.

1. Introduction

Let B be a convex polygonal object having k vertices and edges, free to move (translate and rotate) in an open two-dimensional space V bounded by a collection of polygonal obstacles (“walls”) having altogether n corners. The problem studied in this paper is to plan automatically a continuous obstacle-avoiding motion of B between any two specified initial and final placements; see Fig. 1.1.

This problem was first considered by Schwartz and Sharir [SS], who present an $O(n^5)$ algorithm for its solution (which applies to nonconvex polygonal moving objects as well). Since then this classical “piano-movers” problem had been

* Work on this paper by the second author has been supported by Office of Naval Research Grant N00014-82-K-0381, National Science Foundation Grant No. NSF-DCR-83-20085, and by grants from the Digital Equipment Corporation, and the IBM Corporation.




Fig. 1.1. An instance of the motion-planning problem.

studied extensively, and several efficient algorithms had since been developed for certain special cases of it [BZ], [LS1], [LS2], [KS], [KLPS], [OY], [OSY1], [OSY2], [SiS].

The algorithm developed in this paper is based on a generalization and combination of the algorithms developed by Leven and Sharir [LS1] and by Sifrony and Sharir [SiS] for the case in which B is a line segment. The Leven-Sharir algorithm partitions the (three-dimensional) space FP of free placements of B (also known as the *free configuration space* of B) into simple, openly disjoint and connected cells, and then determines the adjacency between these cells. This yields an abstract representation of FP by a *connectivity graph* whose nodes are these cells and whose edges connect pairs of adjacent cells. Once the cells containing the specified initial and final placements of B are determined, the motion-planning problem is then reduced to a simple graph searching.

The second algorithm [SiS] also reduces the problem to a combinatorial graph searching, but uses a different graph, called the *vertex graph*, whose nodes are the corners of FP , and whose edges connect pairs of corners that are adjacent along edges of FP , or along some additional auxiliary arcs in FP .

Our algorithm constructs an “intermediate” kind of graph, which we call an *edge graph*; its nodes are edges of the boundary of FP and its edges connect pairs of adjacent FP -edges (in a sense to be defined more precisely below). Our algorithm begins (as in [LS1]) by restricting the motion of B to be purely translational at some fixed orientation θ . This motion has only two degrees of freedom, so that it is easier to calculate its associated restricted two-dimensional space FP_θ of free placements (a task which has already been carried out in [KS]; see also [BZ], [KLPS], and [LS2]), and to represent it as the union of two-dimensional polygonal regions having simple shape. Each such region can be given a discrete combinatorial labeling that does not depend continuously on θ . Roughly speaking, we represent FP_θ by a graph VG_θ , whose nodes are the corners of FP_θ , and whose edges connect pairs of adjacent corners. The nodes of VG_θ are given discrete combinatorial labels (that do not depend continuously on θ).

Next we observe that this combinatorial description of FP_θ will not change as θ varies slightly, unless θ is one of finitely many *critical orientations*, at which some critical condition, which affects the combinatorial structure of FP_θ , occurs.

As it turns out, the most complex of these critical orientations are those at

which the object B makes three simultaneous contacts with the obstacles, without penetrating into any obstacle. If B is a line segment (a “ladder”), then it is shown in [LS1] that the total number of such critical placements of B is $O(n^2)$, which consequently leads to an $O(n^2 \log n)$ algorithm for the desired motion planning. If B is a convex k -gon which is free only to translate in V but not to rotate, then the motion-planning problem becomes simpler and can be accomplished in time $O(kn \log kn)$ [LS2], [KS], [KLPS], [BZ]. This follows from the property, which is proved in [KS] and will be used below, that the number of placements of B (all having the same given orientation) at which it simultaneously touches two obstacles, without penetrating into any obstacle, is $O(n)$ (provided that B and the obstacles are in “general position”; see [KS] and below). If B is also allowed to rotate, then the corresponding critical orientations are much harder to analyze. Since each contact of B with the walls is a contact of either a corner of B with a wall edge or an edge of B with a wall corner, a crude and straightforward upper bound on the number of these critical placements of triple contact of B is $O((kn)^3)$. Moreover, if B is nonconvex, then there are cases where the number of these critical placements of B is indeed $\Omega((kn)^3)$. However, a recent result of Leven and Sharir [LS3] shows that if B is convex, then the number of these critical orientations is only $O(kn\lambda_s(kn))$, where $\lambda_s(q)$ is the maximal number of connected portions of the graphs of q continuous functions which compose the graph of their lower envelope, where it is assumed that each pair of these functions intersect in at most s points. It is shown in [Sz] that $\lambda_s(n) = O(n \log^* n)$ (where $\log^* n$ is the length of the smallest exponential tower $2^{2^{\dots^2}}$ exceeding or equal to n). A better asymptotic bound is given in [HS] for the case $s=3$ and in [Sh1] for larger values of s . These better bounds are roughly of the form $O(n\alpha(n)^{O(\alpha(n)^{s-3})})$, where $\alpha(n)$ is the functional inverse of Ackermann’s function, and is extremely slowly growing. In short, for a fixed s , $\lambda_s(n)$ is nearly linear in n , although, as is shown in [HS] and [Sh2] one has $\lambda_s(n) = \Omega(n\alpha^{\lfloor (s-1)/2 \rfloor}(n))$, so that $\lambda_s(n)$ is superlinear in n for $s \geq 3$.

Using these bounds together with the techniques of [LS1], [SiS], [KLPS], and [KS], we next extend each node in VG_θ , by varying θ , into a node which represents an edge of FP , and then construct an *edge graph* EG which represents adjacency of such edges along the boundary of FP . All this finally yields a motion-planning algorithm for a convex polygonal object B which runs in time $O(kn\lambda_s(kn) \log kn)$. This algorithm is being implemented on an IBM robot RS/2, and its experimental results will be described in a forthcoming paper.

This paper is organized as follows. In Section 2 we describe the recursive discrete representation of FP by its associated edge-graph EG , as outlined above. In Section 3 we discuss the algorithmic details involved in the construction of EG , and in its use for actual motion planning.

2. Discrete Recursive Representation of the Free Configuration Space

Let B be a bounded convex k -sided polygonal body, whose interior is nonempty. B is free to move (translate and rotate) in a bounded two-dimensional open

region V having a polygonal boundary with n corners altogether. We assume that this boundary can be partitioned into a collection of pairwise disjoint simple closed polygonal curves (which we call “walls”); the “wall region” V^c can then be partitioned into a collection of convex polygonal regions having pairwise disjoint (and nonempty) interiors. Note that these assumptions exclude degenerate configurations in which a wall is just an isolated corner or an isolated segment, or in which a wall corner is adjacent to more than two wall edges.

Let P be an arbitrary fixed reference point in the interior of B , and let Q be an arbitrary corner of B . Each placement Z of B in the plane can be represented by the parameters (X, θ) , where X is the position of P , and where θ is the orientation of the vector PQ .

As in [SS], we define a *free placement* (X, θ) of B to be a placement at which B is fully contained within V ; a *semifree placement* (X, θ) of B is defined to be a placement at which B may touch some walls, but not penetrate into the interior of the wall region V^c . The set FP of all free placements of B is an open three-dimensional manifold, and the set SFP of all semifree placements is closed.

As in [LS3], we assume that the moving object B and the obstacles are in *general position*. Roughly speaking, this means that (i) there does not exist a placement of B in which it meets four independent constraints involving contacts with obstacles; and (ii) there do not exist two placements of B with the same orientation such that B meets at each of them three independent constraints involving contacts with obstacles (see [LS3] for more detail). We also assume for simplicity that no wall edge is horizontal; this can always be enforced by an appropriate rotation of V .

Since the motion of B has three degrees of freedom, we first analyze, as in [LS1], only purely translational motion of B (involving just two degrees of freedom), and only then treat the case of general motion of B , including rotation. This will enable us to obtain recursively a combinatorial representation of each cross-section of FP at a fixed θ , from which we will then construct a certain discrete graph which represents the entire space FP in a “connectivity preserving” manner, and which allows us to reduce the motion-planning problem to a discrete problem of path searching through that graph. (Our method can be regarded as a hybrid of the two techniques presented in [LS1] and [SiS], in a sense that will become clear later on. An attempt at direct generalization of the technique of [LS1] has led to certain technical difficulties that we have not been able to overcome.)

2.1. The Case of Translational Motion of B

Definition 2.1.

- (a) [LS3] A (potential) *contact pair* O is a pair (W, S) such that either W is a wall edge and S is a corner of B , or W is a wall corner and S is a side of B . In the first case we call the pair a contact pair of type I and in the second case a contact pair of type II.

- (b) A contact pair of type III is a pair $O = (W, S)$ where W is a wall corner and S is a corner of B .
- (c) An actual *obstacle contact* (i.e., a contact of B with an obstacle) is said to *involve* the contact pair $O = (W, S)$ if this contact is of a point on S against a point of W , and furthermore if this contact is *locally free*, i.e., the inner angle of B at S lies entirely on the exterior side of W if S is a corner of B , and the entire angle within the wall region V^c at W lies exterior to B if W is a wall corner.
- (d) (See [KS].) Let A be one of the convex polygonal obstacles into which V^c is decomposed. The *expanded obstacle* A_θ^* associated with A for a given orientation θ of B is the pointwise vector difference $A - B_\theta$, where B_θ is the standard placement of the moving object B in which P lies at the origin, rotated by θ . A_θ^* is also a convex polygonal region whose sides are vector (Minkowski) differences of the form $W - S$ where (W, S) is a contact pair of type I or II, and whose vertices have a similar representation for contact pairs (W, S) of type III.

It follows from the results of [KS] that the restricted free configuration space FP_θ , which is the space of all free placements of B having orientation θ , can be represented as the complement of the union

$$K_\theta = \bigcup_{i=1}^m (A_i)_\theta^* = \bigcup_{i=1}^m (A_i - B_\theta),$$

where A_1, \dots, A_m are the convex polygonal regions into which V^c is decomposed. See Fig. 2.1 for an illustration of K_θ and of FP_θ . The boundary of K_θ (and also of FP_θ) thus consists of a collection of polygonal curves having finitely many corners. An edge of the boundary of K_θ is a connected portion of an edge of an expanded obstacle induced by a type I or a type II contact pair, and each vertex of K_θ is either a (convex) corner of an expanded obstacle induced by a type III contact pair, or a (nonconvex) intersection point of two sides of different expanded obstacles, each induced by a contact pair of type I or of type II.

The nonconvex corners of the boundary of K_θ are the locations of the reference point P at semifree placements in which B (at orientation θ) makes two distinct




Fig. 2.1. K_θ and FP_θ .

obstacle contacts simultaneously (in each of which either a corner of B touches a wall edge or a side of B touches a wall corner).

It is shown in [KS] (see also [KLPS] and [LS2]) that, for each fixed orientation θ of B , the number of nonconvex corners of K_θ is only $O(n)$ (actually it is only $O(m)$). An algorithm for the calculation of FP_θ , which makes use of this property, is presented in [KS]; its time complexity is $O(kn \log^2 kn)$, which has later been improved to $O(kn \log kn)$ in [LS2], using a different approach involving generalized Voronoi diagrams (and also in a recent paper [BZ] that uses a line-triangle representation of the moving object).

The first step of our algorithm is to use the results of [KS] or of [LS2] to obtain a decomposition of FP_θ into connected components, and to represent each such component Q by a discrete connected graph whose nodes are the corners of Q , each of which is given a discrete combinatorial labeling that does not depend continuously on θ . (This differs from the approach used in [LS1], namely to partition FP_θ into trapezoidal cells, and to establish adjacency of these cells in FP_θ , obtaining in this way a connectivity graph CG_θ whose connected components correspond in a 1-1 manner to the connected components of FP_θ . Although this approach can be used in our case as well, it creates technical difficulties when we attempt to extend these cells into three-dimensional cells as we add the degree of freedom of rotation. We therefore prefer to use the approach mentioned above, which is more similar to that of [SiS].)

More precisely, let $C_0 = C_0(\theta)$ be the collection of all convex and nonconvex corners of K_θ (and also of FP_θ). Define a *vertically maximal* corner of FP_θ to be a convex corner u of K_θ which has the largest y -coordinate among all points of K_θ in a sufficient small neighborhood of u . For each such vertically maximal corner u , we introduce an auxiliary corner u^* , which is the unique point of the boundary of FP_θ that lies directly above u and is such that the open segment uu^* is wholly contained in FP_θ . (Note that since V is assumed to be bounded, FP_θ is also bounded, and u^* thus always exists and is well defined.) Let $C^* = C^*(\theta)$ denote the set of all such auxiliary corners.

Next define a *vertex graph* VG_θ , whose set of nodes is $C = C(\theta) = C_0 \cup C^*$, and whose edges either connect pairs of adjacent corners (including auxiliary ones) along edges of FP_θ , or connect vertically maximal corners u to their associated auxiliary corners $u^* \in C^*$.

For certain critical values of θ , the boundary of FP_θ may contain a corner u for which the intersection of FP_θ with any sufficiently small neighborhood of u is disconnected; this would be the case if at orientation θ a corner of B touches a wall corner, while B also makes another contact with the walls, which corresponds to a convex corner of K_θ also lying on another edge of K_θ . (Note that more degenerate double contacts, in which, say, two convex corners of K_θ touch one another, are ruled out by our assumptions that the obstacles are in general position; see also [LS3].) In these cases split u into two distinct corners, one for each connected component of $FP_\theta \cap N$, for an arbitrarily small circular neighborhood N of u , and then connect each of these split corners u' to its two adjacent corners along the two rays bounding the component Q' of FP_θ associated with u' . The split corner u' will be regarded as lying only on the boundary of Q' , and




Fig. 2.2. FP_θ and VG_θ .

not on the boundary of the other components of FP_θ which lie near u . (Note that, assuming general position of the obstacles, such a corner u cannot be vertically maximal.)

In the following subsection we will consider general motion of B including rotation. It is important for analysis of such general motion that VG_θ does not depend continuously on θ , but rather remains invariant except at certain finitely many *critical orientations*, to be defined below. For this reason we assign discrete combinatorial labelings to each node (and thus also to each edge) of VG_θ in the following simple manner: each convex corner of K_θ is labeled by the type III contact pair that induces it; each nonconvex corner is labeled by the two contact pairs (of type I or II) that induce it; and each auxiliary corner u^* is labeled as an auxiliary corner associated with the type III contact pair that induces the vertically maximal corner u to which u^* corresponds. It is easily checked that these labelings uniquely define the corresponding corners of FP_θ (at a given orientation θ). This labeling scheme turns VG_θ into a discrete structure, clearly not depending continuously on θ .

Theorem 2.1. *Two corners $u, v \in C$ belong to the same connected component of VG_θ if and only if they lie on the boundary of the same connected component of FP_θ .*

Proof. If the edge (u, v) belongs to VG_θ then by definition u and v lie on the (boundary of the) same connected component of FP_θ . Thus the “only if” part of the theorem follows by transitive closure. As to the “if” part, it is clearly true in the case in which u and v lie on the same connected component of the boundary of a connected component Q of FP_θ . If Q has more than one boundary component, then we let A_u (resp. A_v) denote the set of all nodes in VG_θ reachable from u (resp. from v) by a path in VG_θ . We claim that both A_u and A_v contain corners lying on the (unique) exterior boundary E of Q , so that the preceding argument implies that $A_u = A_v$, or, in other words, u and v lie in the same connected component of VG_θ . Indeed, suppose that A_u contains no corner of E . Let $z \in A_u$ be the corner with highest y -coordinate. It is then easy to show that z is a vertically maximal corner. (This is because z lies on an interior boundary of Q , and if the upward-directed vertical ray from z did not contain free positions sufficiently near z , then it would have to intersect the boundary component containing z at

a point having larger y -coordinate than z , contradicting our choice of z . This argument is easy in 2-space; a more general argument for higher-dimensional space is given in Section 3 of [SiS].) But then VG_θ contains the edge (z, z^*) , and z^* has larger y -coordinate than z , a contradiction which completes the proof of the theorem. \square

2.2. The Case of a General Motion of B

We now turn to the general case in which B can both translate and rotate. To do so we first consider how the combinatorial characterization of FP_θ , provided by the graph VG_θ , changes as θ varies.

Definition 2.2. An orientation θ of B is called a *critical orientation*, if one of the following conditions occurs:

- (i) There exists a semifree placement of B at orientation θ at which either it makes simultaneously three distinct obstacle contacts involving contact pairs of types I or II, or it makes simultaneously two obstacle contacts, one involving a contact pair of type III and another involving a contact pair of type I or II. In other words, either three edges of expanded obstacles meet at the same nonconvex corner of K_θ , or a convex corner of K_θ meets another edge of K_θ (see Fig. 2.3(a) and (b)).
- (ii) There exists a vertically maximal convex corner u of K_θ whose associated auxiliary corner u^* coincides with a (convex or nonconvex) corner of K_θ (see Fig. 2.3(c)).
- (iii) (a) Two adjacent edges of K_θ become collinear (Fig. 2.3(e)), or
(b) an edge of K_θ becomes horizontal (see Fig. 2.3(d)).

Lemma 2.1. *The vertex graph VG_θ does not change as θ varies in a sufficiently small neighborhood of any noncritical orientation. Furthermore, for each such sufficiently small neighborhood, the connected component $Q(A, \theta)$ of FP_θ , corresponding to a fixed connected component A of VG_θ , varies continuously (in the Hausdorff topology of sets) with θ .*

Proof. First observe that as long as condition (iii)(a) does not arise, each edge or a convex corner of an expanded obstacle, corresponding to some contact pair O , continues to appear on the boundary of that obstacle, and varies continuously with θ . Moreover, the intersection between any two such edges (belonging to different expanded obstacles) is transversal, unless either these edges become collinear, or they meet at an endpoint of one of them. Hence as long as conditions (i) and (iii) do not arise, each such intersection is transversal, and thus also varies continuously with θ . Also, as long as no three edges of expanded obstacles meet at the same nonconvex corner of K_θ , and no convex corner of K_θ meets an edge of that set, it follows that as we vary θ , each nonconvex corner of K_θ , formed by intersection of edges corresponding to two contact pairs O_1, O_2 , will continue




Fig. 2.3

to appear on the boundary of K_θ , continue to be induced by the same two contact pairs, and vary continuously with θ . It follows that as θ varies slightly, no connected subcomponent of FP_θ (or even no component of the boundary of FP_θ) shrinks to a point and disappears, no such component newly appears, no two such components merge into a single component, nor does such a component split into two subcomponents. Furthermore, each connected component of the boundary of FP_θ retains the same combinatorial representation as a circular sequence of convex and nonconvex corners induced by the same combinations of contact pairs, and each of these corners varies continuously with θ . This also implies that each connected component of FP_θ varies continuously with θ .

To complete the proof of the lemma, let u be a vertically maximal corner of K_θ . First note that as long as we do not cross any critical orientation, u remains vertically maximal as θ varies. Also, as long as condition (ii) does not arise, the auxiliary corner u^* associated with u remains on the same edge of K_θ between the same two vertices of VG_θ , and varies there continuously with θ . All these

observations clearly imply that VG_θ remains constant as θ varies through non-critical orientations. \square

Critical orientations of type (i) were analyzed in [LS3], where it was shown that there are at most $O(kn\lambda_\epsilon(kn))$ such orientations. Section 3 presents an algorithm for the calculation of these orientations, which runs in time $O(kn\lambda_\epsilon(kn) \log kn)$.

Critical orientations of type (iii) are trivial to analyze. In fact, there are only $O(kn)$ such orientations, whose calculation is straightforward.

Finally, critical orientations of type (ii) require involved analysis, similar to, but somewhat simpler than that in [LS3]. This analysis will be presented in Section 3 below, where we show that the number of such critical orientations is also at most $O(kn\lambda_\epsilon(kn))$, and that they can also be calculated in time $O(kn\lambda_\epsilon(kn) \log kn)$.

Suppose then that we have calculated all these critical orientations, and sorted them in circular order. Let T denote the sorted collection of all such orientations.

Let θ_0 be a noncritical orientation, and let u be (the discrete labeling of) a corner in VG_{θ_0} . We will associate with u a *lifespan* $L_u = (\theta_1, \theta_2)$, which is the maximal (open) interval containing θ_0 such that $u \in VG_\theta$ for all $\theta \in (\theta_1, \theta_2)$, and such that u does not become coincident with any other corner of VG_θ for any such θ . We also associate a similar lifespan interval $L_e \subset L_u \cap L_v$ with each edge $e = (u, v)$ of VG_{θ_0} , defined in a similar manner. We call the pair (u, L_u) an *extended corner*, and the pair (e, L_e) an *extended edge*.

We next construct an (*extended*) *edge graph* EG , whose nodes are all these extended corners, and each of whose edges connects an extended corner (u, L_u) to another extended corner (v, L_v) , if (u, v) is an edge of VG_θ for some $\theta \in L_u \cap L_v$ (i.e., if the life span of the edge (u, v) is nonempty).

The graph EG is used to represent FP as follows. For each (label of a) corner of some K_θ , the node (u, L_u) of EG represents an edge of the boundary of FP which consists of all placements $Z = (X, \theta)$ of B such that $\theta \in L_u$ and such that at placement Z either B makes a double contact involving the two contact pairs inducing u , if u is a nonconvex corner of K_θ , or B makes an obstacle contact involving the type III contact pair inducing u , if u is a convex corner of K_θ ; if u is an auxiliary corner, then (u, L_u) represents an arc along a face of FP , consisting of placements (X, θ) at which B makes a single obstacle contact induced by some contact pair of type I or II, such that X lies directly above the convex corner of K_θ inducing u .

It is also worth noting that by definition each edge e of FP is the intersection of two surfaces bounding FP , and is thus the locus of placements Z of B at which it satisfies two independent constraints involving contacts with obstacles. Hence at these placements either B makes two distinct contacts with the walls, involving two distinct contact pairs $O_1 = (W_1, S_1)$, $O_2 = (W_2, S_2)$ of type I or II, or it makes a single contact of one of its corners S against a wall corner W , i.e., involving the type III pair $O = (W, S)$. It follows that e must correspond to some edge of EG defined either in terms of the nonconvex corner labeled by O_1 , O_2 , or by the convex corner labeled by O , except in degenerate situations in which




Fig. 2.4

all placements along e have the same orientation θ ; these can arise in certain degenerate situations, including, e.g., the following cases (see Proposition 2.1 of [LS3]):

- (a) W_1, W_2 are both wall corners and $S_1 = S_2$ (see Fig. 2.4(a));
- (b) W_1, W_2 are both wall corners and S_1, S_2 are parallel sides of B ;
- (c) W_1, W_2 are parallel wall edges (see Fig. 2.4(b));
- (d) W_1 is a wall edge and S_2 is a side of B parallel to W_1 .

Nevertheless, we will show that the edges of FP that are represented as nodes of EG are sufficient for adequate representation of FP .

The edges of EG also have natural interpretation in terms of FP . Let $\xi = [(u, L_u), (v, L_v)]$ be an edge of EG , and let e_u, e_v denote the two edges of FP corresponding to (u, L_u) and (v, L_v) , respectively. Then ξ represents adjacency of e_u and e_v along a face f of FP ; more precisely, for each $\theta \in L_{(u,v)}$, the cross section of f at θ contains a straight segment whose endpoints are (labeled by) u and v . An exception is when $v = u^*$; then ξ represents in a similar manner adjacency of e_u and e_{u^*} along an auxiliary “vertical” surface inside FP .

We claim that EG captures the connectivity of FP in the following sense:

Proposition 2.1.

- (a) *To each connected component A of EG there corresponds a unique connected component Q of FP , such that the nodes in A represent precisely those edges of FP which bound Q (and which do not lie in a single θ cross section of FP).*
- (b) *Each connected component of FP has at least one bounding edge that is represented by a node of EG .*

Proof. The proof of (b) is straightforward. Indeed, let Q be a connected component of FP , and let $Z = (X, \theta) \in Q$. The cross section $Q_\theta = Q \cap FP_\theta$ of Q contains X and has thus nonempty interior. Since V is assumed to be bounded, so must be Q_θ . Thus Q_θ has a bounded polygonal boundary, which therefore must contain at least one nonconvex corner u of K_θ (e.g., choose u to be the corner of Q_θ with the largest y -coordinate). But then it is clear that (u, L_u) is a node of EG which represents an edge of Q .

The proof of (a) consists of two parts. In the first part we argue that if (u, L_u) , (v, L_v) are two adjacent nodes of EG , then they represent edges of FP which bound the same connected component. This claim follows immediately from definition of EG . Thus it follows by transitive closure that all the nodes in a connected component of EG represent edges of FP bounding the same connected component of FP .

The second part proves the converse statement, namely that any pair of edges of FP which bound the same connected component of FP , and neither of which is contained in a single θ cross section of FP , are represented by nodes of EG which belong to the same connected component of this graph. The proof of this statement consists of the following steps.

The set of all critical orientations partitions the angular space into disjoint open *noncritical intervals*. By our assumptions that B and the obstacles are in general position, it follows that at each critical orientation only a single criticality occurs. Define a decomposition of FP into disjoint connected *cells* as follows. Let I be a noncritical interval, and let A be a connected component of VG_θ for any (hence all) $\theta \in I$. The cell $c = c(I, A)$ of FP associated with I and A is defined as

$$c(I, A) = \{(X, \theta) : \theta \in I, X \in Q(A, \theta)\},$$

where, as above, $Q(A, \theta)$ is the (unique) connected component of FP_θ whose boundary corners appear in A .

It is clear that the cells defined in this manner are open, connected, and pairwise disjoint, and that the union of their closures covers the entire space FP .

Two cells $c = c(I, A)$, $c' = c(I', A')$ are called *adjacent* if the intervals I, I' have a common endpoint θ^* and there exists a free placement $(X^*, \theta^*) \in FP$ lying in the closures of both cells c, c' .

Now let $Z = (X, \theta)$, $Z' = (X', \theta')$ be two free placements lying the same connected component of FP , and let $p(t)$ be a continuous path in FP connecting these two placements. Without loss of generality we may assume that $p(t)$ crosses between the cells of FP as defined above only finitely often, intersecting each cell boundary transversally, and in fact enters each such cell at most once. Let c_1, \dots, c_i be the sequence of cells traversed by $p(t)$ in this order. It is then clear that each pair of cells c_i, c_{i+1} in this sequence are adjacent in the sense just defined. Conversely, if c, c' are adjacent cells then one can show that there exists a free motion of B from any placement in c to any placement in c' , which is contained in the union of the closures of these two cells and which intersects the boundary between c and c' transversally at a single point. (To show this, one has to move B from a placement in c (or in c') to a free placement (X^*, θ^*) lying in both closures of c and of c' ; to argue that such a motion is always possible without having to enter other cells, we use a straightforward generalization of the argument given in Section 2.2 of [LS1]; in fact, as argued there, the required motion can be achieved by a pure rotation around (X^*, θ^*) .)

We can therefore define a *connectivity graph* CG , whose nodes are the cells in the above decomposition of FP , and whose edges connect pairs of adjacent cells. The above argument then implies that two cells c, c' lie in the same connected

component of CG if and only if they are contained in the same connected component of FP .

To establish our claim it is thus sufficient to prove that:

- (a) The boundary of each cell c of FP contains at least one edge which is contained in an edge e_u represented by some node (u, L_u) of EG .
- (b) If two cells $c = c(I, A)$, $c' = c(I', A')$ of FP are adjacent, and if (u, L_u) , (v, L_v) are two nodes of EG whose corresponding edges e_u, e_v intersect the boundaries of c, c' respectively, then they lie in the same connected component of EG . Moreover, the same property also holds if $c = c'$.

To prove property (a), let $c = c(I, A)$ be such a cell of FP , and let u be a node in A . Then, for any $\theta \in I$, u represents a corner of the connected component $Q = Q(A, \theta)$ of FP_θ . It thus follows by definition that I is a subset of L_u and that e_u has the desired property.

As to property (b), we first claim that, given two adjacent cells $c = c(I, A)$, $c' = c(I', A')$, there exists a node (w, L_w) of EG representing an edge of FP which meets the boundaries of both c and c' . Indeed, let θ^* be the common endpoint of I and I' . By assumption, θ^* involves only one criticality. If this criticality does not affect the components A, A' of VG_{θ^*} , then clearly they must be equal, and any corner w in $A = A'$ yields a node (w, L_w) of EG with the desired property.

Suppose next that at θ^* a criticality which affects A or A' occurs. Let $Q = Q(A, \theta)$ for $\theta \in I$ (resp. $Q' = Q(A', \theta')$ for $\theta' \in I'$) be the connected component of FP_θ (resp. of $FP_{\theta'}$) represented by A (resp. by A'). Note that both Q and Q' are bounded polygonal regions, so that the exterior boundary of each must have at least three corners. Furthermore, each criticality can affect at most two corners of Q (resp. of Q'). Indeed, the only type of criticality that might affect three corners of, say, Q is that in which three edges of K_θ meet at a common point at $\theta = \theta^*$. But this can affect *three* corners of Q only if Q is bounded by just these three edges, and thus shrinks to a point as $\theta \rightarrow \theta^*$; this however is impossible because then c and c' could not have been adjacent in FP .

These arguments are easily seen to imply that A and A' must share a common node w which is not affected by the critical change occurring at θ^* , and it then follows by definition that (w, L_w) meets the boundaries of both c and c' , as asserted. It is thus sufficient to prove the second claim in (b), because the first claim will then follow by combining the facts that (u, L_u) and (w, L_w) (as edges bounding c) lie in the same connected component of EG , and that (v, L_v) and (w, L_w) (as edges bounding c') also lie in the same connected component of EG .

Thus let (u, L_u) , (w, L_w) be two nodes in EG representing edges of FP both meeting the boundary of the same cell $c = c(I, A)$. Then it follows from Theorem 2.1 that u and w belong to the same connected component of VG_θ , for each $\theta \in I$, so that, by definition, (u, L_u) and (w, L_w) belong to the same connected component of EG . The proof of (b), and hence also of part (A) of our proposition, is thus completed. \square

Definition 2.3. Using the observations made in the proof of (B) above, we can also define a map $\Phi(Z)$ which maps each $Z \in FP$ to a node of EG representing

an edge of FP which bounds the connected component of FP containing Z . Specifically, given any free placement $Z = (X, \theta)$ of B , we can translate B upward from Z (i.e., in the positive y direction) until it reaches a new placement $Z^* = (X^*, \theta)$ at which it makes contact with an obstacle. If this contact involves a contact pair of type III, or consists of two simultaneous contacts involving pairs of types I or II, then X^* is a corner of FP_θ , and we let $\Phi(Z)$ be the node of EG induced by that corner. Otherwise, we continue to translate B from Z^* leftward, maintaining the obstacle contact that involves the same contact pair, until we reach a placement $Z^{**} = (X^{**}, \theta)$ for which X^{**} is a corner of FP_θ , and then continue as above. Again, the fact that FP_θ is bounded implies that Z^* , Z^{**} , and thus also $\Phi(Z)$, are always well defined.

We thus obtain the following main theorem of this section:

Theorem 2.2. *Let $Z, Z' \in FP$ be two free placements of B . Then there exists a collision-free motion of B from Z to Z' if and only if the nodes $\Phi(Z), \Phi(Z')$ of EG as defined above belong to the same connected component of EG .*

3. Algorithmic Details

We now turn to describe an efficient implementation of our motion-planning solution. We first present an algorithm that computes VG_θ for any fixed θ ; this is an expansion of the algorithm described in [KS] for the calculation of FP in the case of purely translational motion of a convex polygonal body amidst convex polygonal obstacles. In the next subsection we show, in a manner similar to [LS3], that the number of critical orientations of type (ii) is $O(kn\lambda_6(kn))$. We follow with a detailed description of the algorithms that find all critical orientations, and then construct the extended edge graph EG . Finally, using these structures, we show how to plan continuous obstacle-avoiding motion of B between any specific initial and final placements.

3.1. Constructing VG_θ

Let θ be a noncritical orientation. The following algorithm constructs VG_θ in four stages. The first three stages produce FP_θ in essentially the same manner as in [KS], and the additional fourth stage computes VG_θ from FP_θ .

(a) For each of the convex polygonal obstacles A into which V^c is decomposed (see Subsection 2.1), calculate the expanded obstacle $A_\theta^* = A - B_\theta$ in time linear in the sum of the number of corners of A and of B (using, e.g., the technique of [GRS]; note that the number of corners of A_θ^* is also bounded by that sum [GRS]). Thus the total number of convex corners of all the expanded obstacles is $O(km + n)$ and they can all be calculated in time $O(km + n)$, where m is the number of the convex subparts of V^c , and k, n are as above. Since $m \leq n$ we can denote this complexity as $O(kn)$. However, if $m \ll n$ then $O(kn)$ may be a gross

overestimation of the actual complexity of the expanded obstacles at any given θ . See below for a discussion on how this observation affects the overall complexity of our algorithm.

(b) Calculate the boundary of K_θ , which is the union of all these expanded obstacles. To do this we use a technique due to [OWW] which calculates the boundary of the union of several superimposed polygonal planar regions in time $O((n+t) \log n)$, where n is the number of line segments constituting the boundaries of these regions, and t is the number of intersections between these segments. We combine this technique with the following divide-and-conquer approach:

1. Recursively find

$$G = \bigcup_{i=1}^{\lfloor m/2 \rfloor} (A_i)_\theta^*$$

$$H = \bigcup_{i=\lfloor m/2+1 \rfloor}^m (A_i)_\theta^*.$$

2. Find the contour of $K_\theta = G \cup H$ using the technique of [OWW].

Since each point of intersection of the boundaries of G and H must lie on the boundary of their union, it follows by a theorem proved in [KS], that the number of these intersection points is only $O(m)$, so that step 2 of the above algorithm runs in time $O(kn \log kn)$, and thus the entire algorithm runs in time $O(kn \log kn \log m)$.

(c) Find the connected components of FP_θ and represent them as an “inclusion tree”, each of whose nodes is a connected component of the boundary of FP_θ , as follows [KS]. Note that FP_θ is a general, possibly disconnected polygonal planar region, whose components need not be simply connected. However, since θ is assumed to be noncritical, each connected component of the boundary of FP_θ must be a closed simple polygonal path. In the inclusion tree that we construct each node represents a simple connected component of this boundary, and such a component C_2 is a child of another C_1 if C_2 is contained in C_1 , and the region between them is either contained entirely within FP_θ or is disjoint from FP_θ . The root of the inclusion tree is a nominal “infinite boundary” consisting of all points at infinity. Each connected component C of FP_θ is represented in the inclusion tree by the node which represents the exterior polygonal boundary of C , and by all children of that node (if any), which represent the connected components of the interior boundary of C . The entire tree is easily calculated in time $O(kn \log kn)$, using a straightforward sweeping technique (see [KS] for more details).

(d) Finally we construct VG_θ from FP_θ as follows. Initially, we take all corners of FP_θ to be vertices of VG_θ , and connect them in VG_θ along edges of FP_θ . Then, sweeping a vertical scan line across the plane, we can easily calculate the auxiliary corner associated with each vertically maximal corner of K_θ , in overall time $O(kn \log kn)$. We then add each of these auxiliary corners u^* as a vertex of VG_θ , split the edge e of FP_θ containing u^* into two subsegments at u^* , remove e from VG_θ , and finally connect u^* to the vertically maximal corner u that induced it, and to the two endpoints of e .

Clearly, the time complexity of the entire algorithm is $O(kn \log kn \log m)$.

Lemma 3.1. *The vertex graph VG_θ is planar, has $O(kn)$ edges and corners, and each vertex in VG_θ is incident to at most three edges.*

Proof. The planarity follows immediately from the fact that FP_θ is planar and the added vertical segments which emanate upward from vertically maximal corners do not intersect any other segment except at a vertex of VG_θ . The number of corners and edges of FP_θ is $O(kn)$ and the number of auxiliary corners does not exceed the number $O(kn)$ of convex corners of FP_θ , and each such auxiliary corner adds two edges of VG_θ . Hence VG_θ has $O(kn)$ edges and corners. Finally, it is plain that each vertex in VG_θ is incident to two edges, except for the vertically maximal vertices and their corresponding auxiliary vertices, each being incident to three edges. \square

3.2. The Number of Critical Orientations of Type (ii)

We next turn to the problem of estimating the number of critical orientations of type (ii). Recall that these orientations arise when the auxiliary corner u^* associated with a vertically maximal corner u of K_θ coincides with a convex or a nonconvex corner of K_θ . Also recall that a vertically maximal corner u is induced by a contact pair $O = (W, S)$, where W is a wall corner and S a corner of B such that u has the largest y -coordinate among all points of K_θ in a sufficiently small neighborhood of u .

The number of critical orientations θ at which such u lies immediately below a convex corner ν of K_θ is $O(k^2n^2)$ and a superset of these orientations, also of size $O(k^2n^2)$, can be calculated in a straightforward manner by iterating over all $O(k^2n^2)$ possible two contact pairs of type III. We thus proceed to estimate the number of critical orientations at which u lies directly below a nonconvex corner.

Let u be such a fixed vertically maximal corner induced by a type III contact pair $O = (W, S)$. For each contact pair $O_1 = (W_1, S_1)$ of type I or II, and for each θ define the function $G_{OO_1}(\theta)$ to be the distance along the vertical ray from the corner u upward to its intersection with the edge induced by the pair O_1 if such an intersection exists, and $+\infty$ otherwise.

Let $\Pi = \Pi_{OO_1}$ denote the domain of definition of $G_{OO_1}(\theta)$. Note that Π need not in general be connected, but, as in [LS3], we have:

Lemma 3.2. Π_{OO_1} consists of at most four subintervals.

Proof. Fix the corner S of B at the wall corner W and let B rotate about that common corner. Each point X in B then traces a circular arc, denoted by C_X , about W . Suppose first that S_1 is a corner of B (and that W_1 is a wall edge). It is clear that $G_{OO_1}(\theta)$ is defined in this case if and only if the vertical ray emanating upward from S_1 intersects W_1 , and that $G_{OO_1}(\theta)$ is equal to the length of the vertical segment from S_1 to the intersection point of this ray with W_1 . Similar

considerations apply when S_1 is a side of B (and W_1 is a wall corner). In this case $G_{OO_1}(\theta)$ is defined if and only if the vertical ray emanating downward from W_1 intersects S_1 , and as before $G_{OO_1}(\theta)$ is then equal to the distance between W_1 and the intersection point of that ray with S_1 .

Thus $G_{OO_1}(\theta)$ is defined for orientations θ of B that fulfill the following three conditions:

- (1) The corner u induced by O exists and is vertically maximal.
- (2) B can make at orientation θ an obstacle contact involving the contact pair $O_1 = (W_1, S_1)$.
- (3) The vertical ray emanating upward from u intersects the expanded obstacle edge induced by O_1 .

Thus the domain of $G_{OO_1}(\theta)$ is the intersection of the three domains I_1 , I_2 , I_3 satisfying respectively these three conditions.

A necessary condition for I_1 , the domain of orientations for which u exists and is vertically maximal, to be nonempty is that W itself be vertically maximal, in the sense that it is the point with the largest y -coordinate among all points of V^c in some small neighborhood of W . Assuming that this is the case, it is then clear that the angular interval where u exists and is vertically maximal must be contained in the interval of orientations of B whose corresponding positions of P on the circle C_P lie in the upper half of that circle. It is also clear that the domain of orientations at which a contact induced by O is possible, is an angular interval whose length is equal to the exterior angle at W minus the angle of B at S . Thus I_1 is an angular interval of length less than or equal to π (see Fig. 3.1(a)).

It is also clear that I_2 , the domain of orientations at which the contact O_1 is possible, is an angular interval whose length is equal either to the exterior angle of B at S_1 minus π if S_1 is a vertex of B , or to π minus the convex angle at W_1 , if W_1 is a wall corner. Thus the intersection of I_1 and I_2 is an angular interval of length strictly less than π (see Fig. 3.1(b)).

As to I_3 , suppose first that W_1 is a wall edge. Let l_1 (resp. l_2) be a vertical ray emanating from the rightmost (resp. leftmost) endpoint of W_1 downward, and consider the infinite trapezoidal strip bounded by W_1 , l_1 , and l_2 . It is easily checked that I_3 consists of these orientations of B at which S_1 lies within that trapezoid, or, in other words, each connected component of I_3 corresponds to a connected component of the intersection of C_{S_1} with this trapezoid. Thus in this case I_3 consists of at most three subintervals (see Fig. 3.1(c)). Furthermore, if I_3 does consist of three subintervals, then each of them has length at most π .

Suppose next that W_1 is a corner, and let l_1 be a vertical ray emanating from W_1 downward. Then I_3 consists of these orientations of B at which the side S_1 intersects l_1 . Here two conditions need be met simultaneously. First, W_1 must lie above the line containing S_1 and W_1 must lie on different sides of this line, and second, the two endpoints of S_1 must lie on different sides of l_1 . As to the first condition, it is easily checked that it can be expressed as a linear inequality in $\sin \theta$ and $\cos \theta$ of the form $p \sin(\theta - \alpha) > q$, where p and q are both positive, and can therefore have a solution set consisting of at most one angular interval of length at most π (see Fig. 3.1(d)). The second condition can change from




Fig. 3.1

being true to being false or vice versa whenever one of the endpoints of S_1 crosses the line containing I_1 . Thus there are at most four orientations at which the second condition can change, again yielding at most two angular intervals at which this second condition can hold. Thus, intersecting the first interval with the latter two, we conclude that I_3 consists of at most three angular intervals in this case too, where each of these subintervals is of length at most π .

It follows that both cases $I_1 \cap I_2 \cap I_3$ consists of at most three angular intervals. Transforming these intervals to the range $0 \leq \theta \leq 2\pi$, we conclude that the domain Π_{OO_i} of $G_{OO_i}(\theta)$ consists of at most four subintervals in this linearized range. This completes the proof of the lemma. \square

If Π_{OO_i} is indeed not connected we will consider each connected portion of $G_{OO_i}(\theta)$ as a separate partially defined function and will still have at most $O(kn)$ such functions for each fixed contact pair O .

We next define the lower envelope of the set of functions G_{OO_i} for O :

$$\Psi_O(\theta) = \min_i G_{OO_i}(\theta).$$

Note that if I is a (maximal) interval over which Ψ_O is attained by a single function G_{OO_i} , then for each $\theta \in I$ the auxiliary corner u^* lies on the expanded obstacle edge induced by O_i . Furthermore, each breakpoint θ on the envelope $\Psi_O(\theta)$ (i.e., a point at which $\Psi_O(\theta)$ is attained simultaneously by two functions G_{OO_i}, G_{OO_j}) is an orientation in which the edge of K_θ containing u^* changes

(from the one labeled by O_i to that labeled by O_j); conversely, every such change in the edge containing u^* corresponds to a breakpoint of $\Psi_O(\theta)$.

Next we show that each pair of the functions G_{OO_i} intersect in at most four points:

Lemma 3.3. *Let $O = (W, S)$ be a contact pair inducing a vertically maximal corner u and let O_1, O_2 be two other distinct contact pairs. Then there are at most four orientations θ in which u lies directly below the nonconvex corner formed by the intersection of the expanded obstacle edges induced by O_1 and O_2 , respectively.*

Proof. Follows immediately from the analysis in [SS] and Lemma 2.2 in [LS3]. Specifically, consider the curve γ_{O_1, O_2} traced by S as B makes simultaneously two obstacle contacts involving the pairs O_1 and O_2 . As shown in [SS], this curve is either a straight segment or part of an ellipse or part of a quartic algebraic curve. The desired placement at which u lies directly below the intersection point of the edges induced by O_1 and O_2 is the intersection between this curve and the vertical line emanating upward from W . Since γ_{O_1, O_2} is at most quartic there are at most four such intersections. \square

It therefore follows (see [At]) that the number of breakpoints along Ψ is at most $\lambda_6(kn)$. To see this, it suffices to show that each pair $g_i = G_{OO_i}, g_j = G_{OO_j}$ of these partial functions can alternate along the lower envelope Ψ_O at most seven times. Indeed, the two extreme portions in such an alternation can occur within the symmetric difference of the domains of g_i and of g_j (observe that this symmetric difference consists of at most two intervals, on each of which just one of these functions is defined), and at most five additional alternating portions of the graphs of g_i, g_j can occur along the lower envelope within the intersection of the domains of these two functions, because g_i and g_j must intersect between each pair of such adjacent alternating occurrences, and by Lemma 3.3 only four such intersections can occur. Since the number of vertically maximal corners is at most $O(kn)$, we have in summary.

Theorem 3.1. *The total number of critical orientations of type (ii) is at most $O(kn\lambda_6(kn))$.*

An efficient procedure for calculating these orientations will be described in the following subsections.

3.3. Calculation of All Critical Orientations

We now proceed to describe a collection of algorithms that actually calculate all critical orientations. For each wall edge W we denote its two endpoints as W_L and W_R , so that the wall region lies to the left of the oriented segment $W_L W_R$. Similarly, for each side S of B , we denote its endpoints by S_L and S_R , so that B lies to the right of the oriented segment $S_L S_R$.

Calculation of Type (i) Critical Orientations. This phase extends and provides more detail of the procedure briefly outlined in [LS3]. We recall some definitions and lemmas from [LS3] to be used in the description of this phase.

Definition 3.1 [LS3].

- (a) The *tangent line* T of a contact pair $O = (W, S)$ of type I or of type II is either the line passing through W if W is a wall edge or the line passing through W and parallel to S if W is a wall corner (note that in the second case T depends on the orientation of B).
- (b) Let O_1, O_2 be two contact pairs for which there exists a (unique) placement $Z = (X, \theta)$ of B at which it makes simultaneously two obstacle contacts involving O_1 and O_2 , respectively. We say that O_2 *bounds* O_1 at the orientation θ if $B^* = \text{conv}(S_1 \cup S_2)$ always intersects W_2 as we translate B from Z along the tangent T_1 in the direction of the intersection point z_{12} of T_1 and T_2 , until the last placement at which S_1 still touches W_1 .
- (c) Let $O_1 = (W_1, S_1)$ be a contact pair. If W_1 is a wall edge (and S_1 is a corner of B), the two contacts of S_1 against the two endpoints of W_1 will be called *endpoints* of (the edge induced by) O_1 . If W_1 is a wall corner, we define the endpoints of (the edge induced by) O_1 in a symmetric manner as contacts of the endpoints of S_1 against W_1 .
- (d) We will say that at orientation θ , O_2 bounds O_1 toward an endpoint E of the edge induced by O_1 , if the translation of B from Z as in (b) above terminates at the contact E .

Proposition 3.1 [LS3]. *If at some placement $Z = (X, \theta)$, B makes simultaneously two obstacle contacts involving the contact pairs O_1, O_2 , then either O_1 bounds O_2 at θ , or O_2 bounds O_1 at θ , except in the degenerate case where the tangents T_1, T_2 are coincident or parallel.*

Let O_1 be any contact pair and let E be one of its endpoints. Consider all contact pairs that bound O_1 (at any orientation θ) toward E . For each such contact pair O_2 we define a *bounding function* $F_{O_1 O_2}(\theta)$, over the domain of orientations θ of B in which O_2 bounds O_1 toward E , to be the distance of the endpoint of W_1 at which the contact E is made from the contact point of S_1 and W_1 , if O_1 is a contact pair of type I, or the distance of the endpoint of S_1 involved in E from the contact point of S_1 and W_1 , if O_1 is a contact pair of type II.

We partition the collection of bounding functions $\{F_{O_1 O_i}\}$ of O_1 into two classes $\mathbf{A}_L(O_1), \mathbf{A}_R(O_1)$ so that for all functions $F_{O_1 O_i}$ in \mathbf{A}_L , O_i bounds O_1 toward the endpoint W_L of W_1 (or S_L of S_1) whereas for all functions in \mathbf{A}_R , O_i bounds O_1 toward the other endpoint of W_1 (or of S_1). Note that there may exist a contact pair O_i for which a bounding function $F_{O_1 O_i}$ appears in both collections $\mathbf{A}_L, \mathbf{A}_R$, but then these two functions will have openly disjoint domains of definition [LS3]. Also the domain of definition of each $F_{O_1 O_i}$ consists of at most five connected intervals [LS3]. If the domain of such a $F_{O_1 O_i}$ is not connected, we regard this function as several (≤ 5) distinct partially defined functions having connected (and pairwise openly disjoint) domains (see [LS3]).

We are now in position to describe the collection of procedures which compute all the type (i) critical orientations.

Step 1: Find All Bounding Functions. For every two contact pairs $O_i = (W_i, S_i)$, $O_j = (W_j, S_j)$, each being either of type I or of type II, find the range of orientations (which can be empty) in which O_i bounds O_j toward a specific endpoint E of O_j . Split the resulting bounding function F_{O_i, O_j} into (at most five) "subfunctions", and add these functions, together with additional information concerning their parameters and domains of definition, to the appropriate collection $A_L(O_j)$ or $A_R(O_j)$. (As mentioned above, for certain pairs O_i, O_j of contact pairs, functions F_{O_i, O_j} may be added to both $A_L(O_j), A_R(O_j)$.)

The results of [LS3] reviewed above clearly imply that the storage required up to this point does not exceed $O(k^2 n^2)$. Each bounding function (for two fixed contact pairs) can be processed in $O(1)$ time, and thus Step 1 can be performed in time $O(k^2 n^2)$.

Our next step is to find, for each of the $O(kn)$ contacts pairs O , the lower envelope

$$\Psi_{L;O}(\theta) = \min_i \{F_{OO_i}(\theta) : F_{OO_i} \in A_L(O)\},$$

and the corresponding envelope $\Psi_{R;O}$ for the collection $A_R(O)$.

It was shown in [LS3] that each type (i) critical orientation θ is one of the three following kinds:

- (1) θ is the orientation of an intersection point of two functions F_{OO_i} and F_{OO_j} lying along some lower envelope $\Psi_{E;O}$.
- (2) θ is the orientation of an intersection point of some function F_{OO_i} on $\Psi_{L;O}$ and a transformation of a function F_{OO_j} lying along $\Psi_{R;O}$ (F_{OO_j} is shifted and reflected so that both functions measure distance from the same endpoint of O).
- (3) There exist three contact pairs O_1, O_2, O_3 , such that O_1 bounds O_2 (toward an endpoint E_2), O_2 bounds O_3 (toward an endpoint E_3), and O_3 bounds O_1 (toward an endpoint E_1) at θ , and the envelopes $\Psi_{E_1;O_1}, \Psi_{E_2;O_2}, \Psi_{E_3;O_3}$ are attained at θ by the functions $F_{O_3, O_1}, F_{O_1, O_2}, F_{O_2, O_3}$, respectively.

Step 2: Calculate Lower Envelopes. The following algorithm for calculating the lower envelopes is adapted from [At]. Since both kinds of envelopes are calculated in much the same way, we will describe below only the calculation of the lower envelopes $\Psi_{L;O}$ for the collections $A_L(O)$.

- (1) Fix a contact pair O , and partition $A_L(O)$ into two disjoint subsets A'_L, A''_L of roughly equal size.
- (2) Compute recursively the two lower envelopes

$$\Psi'(\theta) = \min\{F_{OO_i}(\theta) : F_{OO_i} \in A'_L\},$$

$$\Psi''(\theta) = \min\{F_{OO_i}(\theta) : F_{OO_i} \in A''_L\}.$$

Each of these recursive calculations will have produced a sequence of angular intervals in each of which the corresponding partial envelope is

attained by a single bounding function; we refer to the endpoints of these intervals as *breakpoints* along the corresponding envelope.

- (3) Merge these two sequences of intervals to obtain a refined sequence Γ of angular intervals. For each refined interval $I \in \Gamma$ there exist unique contact pairs O', O'' with $F_{OO'} \in \mathbf{A}'_L, F_{OO''} \in \mathbf{A}''_L$, such that $\Psi'(\theta) = F_{OO'}(\theta), \Psi''(\theta) = F_{OO''}(\theta)$ for each $\theta \in I$. By the analysis of [LS3] the two functions $F_{OO'}$ and $F_{OO''}$ intersect in at most four points (some of which may not belong to I), which can be calculated as the roots of some quartic polynomial. Each of these intersections which lies in I is clearly a breakpoint of $\Psi = \Psi_{L,O}$, and every breakpoint of Ψ is either of this kind, or is a breakpoint of Ψ' or of Ψ'' . Thus we can calculate Ψ from the merged sequence Γ in time proportional to the number of intervals in Γ , which, by [LS3], is $O(\lambda_6(kn))$. Thus the total time required by Step 2, when repeated for each contact pair O and extended to calculate also lower envelopes of the form $\Psi_{R,O}$, is $O(kn\lambda_6(kn) \log kn)$.

Remark. At the bottom level of the recursion, when the collection of bounding functions consists of just a single function, the sequence of breakpoints that has to be generated simply consists of the two endpoints of the domain of definition of that function. Thus, some of the breakpoints in the final envelope Ψ will represent endpoints of such domains.

The output of this step consists of $O(kn)$ sorted lists of critical orientations, where each list represents the breakpoints along a single lower envelope $\Psi_{L,O}$ or $\Psi_{R,O}$. We denote by $\Phi_L(O)$ (resp. $\Phi_R(O)$) the list representing $\Psi_{L,O}$ (resp. $\Psi_{R,O}$) for a contact O . For each such list $\Phi_E(O)$, the procedure also produces an associated list $M_E(O)$ of the bounding functions attaining the corresponding envelope on the sequence of intervals delimited by the orientations in $\Phi_E(O)$.

By now we have found (a superset of) all critical orientations θ of type (i) in which B can make simultaneously (at a free placement) obstacle contacts involving three distinct contact pairs O_1, O_2, O_3 , such that two of these pairs, say O_2 and O_3 , bound the third one O_1 toward the same endpoint of O_1 , i.e., such that the bounding functions F_{O_1,O_2} and F_{O_1,O_3} belong to the same collection $\mathbf{A}_L(O_1)$ or $\mathbf{A}_R(O_1)$ and meet one another at θ along the corresponding lower envelope. (Note that not all orientations corresponding to envelope breakpoints necessarily yield free critical contacts of B . However, subsequent steps of the algorithm will be able to detect such spurious breakpoints and discard them; see below for more details.)

Step 3: Calculate the Second Kind of Type (i) Critical Orientations. These are orientations θ at which B makes simultaneously, at some free placement, obstacle contacts involving three distinct contact pairs O_1, O_2, O_3 , such that two of them, say O_2 and O_3 , bound O_1 , but with F_{O_1,O_2} belonging to $\mathbf{A}_L(O_1)$ and F_{O_1,O_3} belonging to $\mathbf{A}_R(O_1)$ at θ . Calculation of such orientations is easy and proceeds as follows:

- (1) For each contact pair O , merge the two lists $\Phi_L(O)$ and $\Phi_R(O)$ of breakpoints to obtain a refined sorted list of angular intervals delimited by these breakpoints.

- (2) For each such refined interval I , there exist unique functions $F_{OO_1} \in \mathbf{A}_L(O)$, $F_{OO_2} \in \mathbf{A}_R(O)$ which attain $\Psi_{L;O}$ and $\Psi_{R;O}$, respectively, over I . Calculate the (at most four) intersections between these two functions and add those intersections that lie in I to the output list of new critical orientations. (Note that in this calculation F_{OO_2} has to be shifted and reflected into the coordinate frame in which F_{OO_1} is defined, so that both functions will measure the distance from the same endpoint of O .)

Clearly, Step 3 runs in $O(kn\lambda_6(kn))$ time.

Step 4: Calculate the Third Kind of Type (i) Critical Orientations. Finally, we calculate the third and most complex kind of critical orientations. At each such orientation θ , B can make simultaneously a free triple contact involving three distinct contact pairs O_1, O_2, O_3 , such that $F_{O_1O_2} \in \mathbf{A}_{E_1}(O_1)$, $F_{O_2O_3} \in \mathbf{A}_{E_2}(O_2)$, $F_{O_3O_1} \in \mathbf{A}_{E_3}(O_3)$, where $E_i \in \{L, R\}$ for $i = 1, 2, 3$ (and such that all these functions lie on the corresponding lower envelopes). In other words, no two of these contact pairs simultaneously bound the third one.

To find these orientations we first merge all lists $\Phi_L(O), \Phi_R(O)$, for all contact pairs O , to obtain a single sorted list Φ consisting of $O(kn\lambda_6(kn))$ refined noncritical intervals delimited by all these breakpoints. Each interval I in Φ has the property that each envelope $\Psi_{E;O}$ is attained over I by a single bounding function $F_{OO'}$.

Step 4 processes the intervals in Φ in order, maintaining a priority queue Q of potential critical orientations of the kind we seek, and a map Ψ^* which maps each pair (O, E) of a contact pair O and endpoint E to the contact pair O' for which $\Psi_{E;O} = F_{OO'}$ over the currently processed interval of Φ . This procedure runs as follows:

- (1) Starting at the initial interval I_0 , we first initialize the value of Ψ^* for I_0 , and then construct a list Λ of potential triple contacts of the third kind induced by the various lower envelopes over I_0 , as follows. For each contact pair O_1 , and each $E_1 \in \{L, R\}$, find the (unique) contact pair O_2 such that $\Psi_{E_1;O_1} = F_{O_1O_2}$ over I_0 (using the map Ψ^*). For each $E_2 \in \{L, R\}$ find the (unique) contact pair O_3 such that $\Psi_{E_2;O_2} = F_{O_2O_3}$ over I_0 . Finally, for each $E_3 \in \{L, R\}$ for which $\Psi_{E_3;O_3} = F_{O_3O_1}$ over I_0 , add the item $(O_1, O_2, O_3, E_1, E_2, E_3)$ to Λ . Clearly, Λ consists of at most $O(kn)$ items.
- (2) Next initialize the priority queue Q as follows. For each item $(O_1, O_2, O_3, E_1, E_2, E_3)$ in Λ , calculate the (at most four) orientations at which a triple contact involving the pairs O_1, O_2, O_3 can be made such that at this contact O_2 bounds O_1 toward E_1 , O_3 bounds O_2 toward E_2 , and O_1 bounds O_3 toward E_3 . Orientations of this kind that belong to I_0 are immediately added to the output list of critical orientations, whereas all the others are added to Q , each tagged by the corresponding item (O_1, \dots, E_3) .
- (3) Next we process the remaining intervals in Φ in order. Let I be the interval presently being processed, let I' be the preceding interval, and let θ be the common endpoint of I' and I . By our assumptions at most one triple contact can occur at θ . Suppose this contact involves three individual

contact pairs O_1, O_2, O_3 such that, say, O_2 and O_3 both bound O_1 toward its L endpoint. Then the only envelope for which the bounding function attaining it changes from I' to I is Ψ_{L,O_1} . Thus if we were to calculate (in the manner described above) a list Λ of potential triple contacts from the envelopes over I , we would have obtained the same list that would be obtained from the envelopes over I' , with the exception of contacts involving O_1 (more precisely, contacts corresponding to items of the form $(O_1, O', O'', L, E', E'')$). But there are clearly only $O(1)$ items of this form, and they can all be computed in $O(1)$ time using the map Ψ^* (for the interval I) as in (1) above. We therefore proceed as follows:

- (i) Update the value of $\Psi^*(O_1, L)$, obtain the $O(1)$ new potential triple contacts arising from this change, and insert the corresponding critical orientations into Q . (Optionally, we may also choose to delete the $O(1)$ entries from Q that were obtained using the old value of $\Psi^*(O_1, L)$.)
- (ii) Remove from Q all entries whose orientations lie in I . Let $(\bar{\theta}, (O, O', O'', E, E', E''))$ be an entry removed from Q . We check whether this contact is still valid, in the sense that the present value of Ψ^* satisfies $\Psi^*(O, E) = O'$, $\Psi^*(O', E') = O''$, and $\Psi^*(O'', E'') = O$. If so, add this entry to the output list.

This concludes the description of Step 4. Its time complexity is clearly $O(kn\lambda_6(kn) \log kn)$. Therefore all critical orientations of type (i) can be calculated in time $O(kn\lambda_6(kn) \log kn)$.

Step 5: Calculation of Type (ii) Critical Orientations. These orientations are calculated in a similar (and simpler) manner to the calculation of type (i) orientations. Recall that a type (ii) orientation θ is one in which a vertically maximal corner u of K_θ induced by a type III contact pair $O = (W, S)$ lies immediately below another (convex or nonconvex) corner u^* of that set. If u^* is convex, then u^* is induced by another contact pair (W_1, S_1) of a wall corner W_1 and a corner S_1 of B . In this case θ can be found by intersecting the circle at radius $|SS_1|$ about W with the vertical ray descending from W_1 . There are at most two such intersections, and for each of these intersection points Z we obtain a corresponding critical orientation θ , at which the segment SS_1 on B becomes parallel to WZ . It is therefore clear that the total number of such critical orientations is at most $O(k^2n^2)$, and that, by repeating the above procedure for every two type III contact pairs $(W, S), (W_1, S_1)$, we obtain in time $O(k^2n^2)$ a superset of the set containing all these critical orientations.

If the corner u^* is nonconvex, then we have to apply the techniques described in Subsection 3.3 to calculate the corresponding critical orientations. Specifically:

- (1) For each type III contact pair $O = (W, S)$ and for each contact pair $O_i = (W_i, S_i)$ of type I or II, find (in constant time) the domain of definition and other parameters of the function $G_{OO_i}(\theta)$. (Recall that this function measures, for each orientation θ , the vertical distance between the corner

- u of K_θ induced by O and the edge e of that set induced by O_i , whenever this distance is well defined and u lies below e .)
- (2) Calculate the lower envelope Ψ_O of the functions G_{OO_i} , over all contact pairs O_i , using the procedure as applied in Step 2 to the functions F_{OO_i} . This produces a list T_O of critical orientations, each representing a break-point along Ψ_O . By the results in Subsection 3.2, T_O consists of at most $O(\lambda_6(kn))$ orientations.
 - (3) Finally, we can merge T_O with the previously calculated list of potentially critical orientations θ at which u lies below another convex corner of K_θ . We can determine whether any such latter orientation θ is indeed critical by testing whether it represents a point actually lying on the envelope Ψ_O , and discard θ if it is found not be critical.

Clearly, Step 5 of our algorithm requires $O(kn\lambda_6(kn) \log kn)$ time.

Step 6: Calculation of Type (iii) Critical Orientations. Recall that this type of critical orientation arises when two adjacent edges of K_θ become collinear, or when an edge of K_θ becomes horizontal. As noted in Subsection 2.2, there are at most $O(kn)$ orientations of the first kind that are not also type (i) critical orientations, and, assuming general position of the walls (so that no wall edge is horizontal), there are at most $O(k)$ orientations of the second kind. All these orientations are trivial to calculate in time $O(kn)$ (note that as before this calculation produces a superset of the required critical orientations, because not all these orientations necessarily arise in semifree critical contacts of B).

3.4. Constructing the Edge-Graph EG

By now we have calculated a set T of size $O(kn\lambda_6(kn))$ that contains all the critical orientations of all three types. Suppose that T is sorted in ascending order, and assume without loss of generality that $\theta = 0$ is not an orientation in T .

As in Section 2, we represent each node ξ of the graph EG by a pair (u, L_u) , where $L_u = (\theta_1, \theta_2)$ is the angular lifespan of ξ , and where u is the (discrete labeling of the) corner of VG_θ , for $\theta \in L_u$, that lies on the edge of FP represented by ξ .

The calculation of EG will be accomplished in a manner similar to that of [LS1]. That is, we process critical orientations in increasing order, maintaining the ‘‘cross-section’’ graph VG_θ and use it to update EG at each critical orientation. At each such orientation θ we determine those nodes of EG whose lifespan terminates at θ (these are nodes whose corresponding corners have to be deleted from VG_θ), and the new nodes whose lifespan starts at θ . Nodes of the first kind will already have been stored in EG , and we update their lifespan by adding θ as its terminal orientation. Nodes of the second kind are added to EG , with θ as the initial orientation of their lifespan, and with the corresponding terminal orientation being presently left undefined. EG is also augmented by edges connecting the nodes just inserted with nodes already present in EG , as described

in Subsection 2.2 (note that these edges correspond to new edges in $VG_{\theta'}$, for θ' slightly larger than θ).

To initialize EG we first calculate the graph $VG_{\theta=0}$. For each corner u of this graph we create a corresponding initial node in EG whose label contains u , but with both initial and terminal lifespan orientations left undefined.

Our algorithm has thus the following structure:

- (1) Merge all the various lists of critical orientations into a single sorted list T consisting of $0 < \theta_1 < \theta_2 < \dots < \theta_w < 2\pi$ (where $w = O(kn\lambda_6(kn))$). (Without loss of generality we assume that 0 is not a critical orientation.)
- (2) Compute VG_θ for $\theta=0$. Copy VG_θ into EG , expanding the label of each corner and edge of VG_θ by two (presently undefined) entries for the two delimiting orientations of the lifespan of the resulting node or edge of EG .
- (3) For each critical orientation θ_i , $i = 1, \dots, w$, do:
 - (a) Determine the changes that VG_θ undergoes as θ crosses θ_i from smaller to larger orientations.
 - (b) Let v_1, \dots, v_l be the corners and e_1, \dots, e_m be the edges of VG_θ that disappear as θ crosses θ_i . Remove them from VG_θ , and assign θ_i as the terminal lifespan orientation of the corresponding nodes and edges in EG .
 - (c) Let v_1^*, \dots, v_s^* be the corners newly appearing in VG_θ after θ crosses θ_i , and let e_1^*, \dots, e_r^* be the new edges in that graph. Add these elements to VG_θ , and add the corresponding nodes and edges to EG with θ_i as their initial lifespan orientation (and with their terminal lifespan orientation presently left undefined).
- (4) {At the end of the above main loop the nodes and edges of EG whose terminal lifespan orientation is still undefined are in 1-1 correspondence with the initial nodes and edges of EG , as created in step (2) above from the initial graph $VG_{\theta=0}$.}

This final step in the construction of EG simply identifies these two types of nodes. Specifically, for each node $(u, (\theta_1, \Omega))$ of EG whose terminal lifespan orientation is still undefined, find the corresponding node $(u, (\Omega, \theta_2))$ of EG whose initial lifespan orientation is undefined and merge them into the single node $(u, (\theta_1, \theta_2))$. (Note that there are only $O(kn)$ such pairs of nodes.)

Remark. In step (3)(a) we can first check whether the criticality that is supposed to happen at θ actually arises in VG_θ ; otherwise θ is a spurious critical orientation which we can simply discard. (For example, if θ is an orientation at which B makes three obstacle contacts simultaneously, then, if this is really a free triple contact then for θ' near θ the boundary of $K_{\theta'}$ must contain edges induced by all these three contacts, and we can easily check whether this is the case from the current value of VG_θ .)

We can now state our main result.

Theorem 3.2. *The graph EG has at most $O(kn\lambda_6(kn))$ nodes and edges. The algorithm just described computes EG in $O(kn\lambda_6(kn) \log kn)$ time.*

Proof. Let θ_0 be a critical orientation in T , and let $(\theta_0 - \epsilon, \theta_0 + \epsilon)$ be a sufficiently small interval in which only θ_0 is critical. As in Section 2, we also make the simplifying assumption that only one critical contact of B occurs at θ_0 . (See also a discussion on this assumption in [LS3]. An exception are orientations of type (iii)(b), at each of which up to $O(n)$ different expanded obstacle edges become simultaneously horizontal; see below for the special treatment of these orientations.) The changes that occur in VG_θ (and induce similar changes in EG) as θ varies from $\theta_0 - \epsilon$ to $\theta_0 + \epsilon$ depend on the type of criticality that arises at θ_0 :

(a) Suppose first that θ_0 is a critical orientation of type (i) at which three edges of K_θ become concurrent at some nonconvex corner of K_θ . Then VG_θ can change in one of the following ways: either

- (i) one corner of this graph disappears and is replaced by two new corners connected by a new edge (see Fig. 3.2(a)); or
- (ii) two adjacent corners of VG_θ merge and form a new corner, and the edge connecting them disappears (see Fig. 3.2(a)); or
- (iii) a small connected component of VG_θ consisting of just three corners shrinks to a single isolated corner and then disappears (see Fig. 3.2(b)); or, finally,
- (iv) a new connected component of VG_θ newly appears, first as an isolated corner and then expands to a triangle (see Fig. 3.2(b)).

Since, by the arguments in Subsection 3.1, the degree of each vertex of VG_θ at any noncritical θ is at most three, it follows that in each of these cases step (3) of the algorithm newly generates or modifies only a small fixed number of nodes and edges of EG as a result of these changes in VG_θ at $\theta = \theta_0$.




Fig. 3.2




Fig. 3.3

(b) Next, suppose that θ_0 is a critical orientation of type (i) at which a convex corner of K_θ comes to lie on an edge of K_θ . In this case VG_θ can change in one of the following ways: either

- (i) one convex corner of that graph disappears and is replaced by two new nonconvex corners; these corners lie in the interior of a former edge of VG_θ , and now split that edge into two new edges of that graph, each delimited by one of these new corners (see Fig. 3.3(a)); or, conversely,
- (ii) two nonconvex corners u, v of VG_θ merge into a new convex corner w ; furthermore, w becomes adjacent to one neighbor u' of u and to one neighbor v' or v , and the two remaining neighbors s, t (of u and v resp.) become adjacent to one another (see Fig. 3.3(a)).

Thus in each of these cases EG undergoes only $O(1)$ changes at θ_0 .

Since there are at most $O(kn\lambda_6(kn))$ critical orientations of type (i), the total number of new nodes and edges added to EG at these orientations is also bounded by $O(kn\lambda_6(kn))$.

(c) Next, suppose that θ_0 is a critical orientation of type (ii), in which a vertically maximal corner u of K_θ comes to lie directly below another corner v of K_θ (i.e., the associated auxiliary corner u^* coincides with v). In this case u^* changes its location from one edge (or portion of an edge) vw_1 of K_θ to an

adjacent one vw_2 . Consequently the two former edges vu^* , u^*w_1 of VG_θ have to be merged into a single edge vw_1 , whereas the edge vw_2 is split into two new edges vu^* , u^*w_2 . (See Fig. 3.3(b); note that the new edge vu^* is a portion of an expanded obstacle edge induced by a contact pair different from that inducing the previous edge vu^* .) Again only $O(1)$ changes have to be made at θ_0 in the graphs VG_θ and EG .

(d) Finally, suppose θ_0 is a critical orientation of type (iii). Suppose first that θ_0 is of subtype (iii)(a), i.e., a side S of B at orientation θ_0 becomes parallel to a wall edge W . Let W_L, W_R be the endpoints of W and let S_L, S_R be the endpoints of S . Without loss of generality suppose that for θ slightly less than θ_0 , K_θ has two adjacent edges induced respectively by the contact pairs (W_L, S) , (W, S_L) , and that these edges become collinear at θ_0 , and, for θ slightly larger than θ_0 , are replaced by two other edges induced respectively by the pairs (W, S_R) , (W_R, S) .

Consequently, for arbitrarily small $\varepsilon > 0$, all the vertices and edges of $VG_{\theta-\varepsilon}$ that include (W_L, S) or (W, S_L) in their label terminate at θ , and corresponding new vertices and edges that include (W, S_R) , (W_R, S) in their label show up in $VG_{\theta+\varepsilon}$. The rest of the vertices and edges of both VG graphs remain unchanged. Since there are at most $O(kn)$ edges and vertices including these pairs in their label in $VG_{\theta-\varepsilon}$ (resp. $VG_{\theta+\varepsilon}$), and the degree of each of these vertices is at most three, the number of edges and vertices removed from (resp. added to) VG_θ and EG at each critical orientation of type (iii)(a) is at most $O(kn)$. Since there are at most $O(kn)$ critical orientations of this type, the total number of modifications of VG_θ and of EG at these orientations is at most $O(k^2n^2)$.

Next, suppose that θ_0 is of subtype (iii)(b), i.e., a side S of B becomes horizontal (with B lying above S) at θ_0 . Let S_L, S_R be the endpoints of S . Then without loss of generality we can assume that for θ slightly less (resp. larger) than θ_0 , S_L is lower (resp. higher) than S_R . Then for θ slightly less than θ_0 , K_θ contains at most $O(n)$ vertically maximal corners induced by contact pairs of the form (W, S_L) , which cease to be vertically maximal as θ_0 is crossed, and have to be replaced by another set of at most $O(n)$ adjacent corners induced by corresponding pairs of the form (W, S_R) which newly become vertically maximal. Thus, arguing as above, at most $O(n)$ changes are required in the graphs VG_θ and EG at each of the $O(k)$ critical orientations of type (iii)(b), thus the total number of graph modifications at these orientations is only $O(kn)$.

This completes the proof of the first part of Theorem 3.2.

Analysis of the Time Complexity of the Algorithm. We have already shown that the preliminary stages of our procedure, which calculate all critical orientations, require $O(kn\lambda_6(kn) \log kn)$ time. As to the complexity of the main part of the algorithm described in this subsection we argue as follows.

In step (1) we merge $O(kn)$ sequences of critical orientations, each having length at most $O(\lambda_6(kn))$; this is easily accomplished in time $O(kn\lambda_6(kn) \log kn)$.

Construction of $VG_{\theta=0}$ in step (2) can be accomplished, e.g., in time $O(kn \log^2 kn)$ using the algorithm described by [KS].

As to step (3), it processes $O(kn\lambda_6(kn))$ critical orientations. At each orientation

θ_i of type (i) or (ii) it has to perform only $O(1)$ updating operations on the graphs VG_θ and EG , each of which can be accomplished in time $O(\log kn)$.

If θ_i is a critical orientation of type (iii)(a), let e_1, e_2 be the two adjacent expanded obstacle edges that become collinear at θ_i , and are then replaced by two other edges, e_3, e_4 (see the paragraph (d) in this proof). Then it is easily checked that the boundary of FP_θ varies continuously with θ (in the Hausdorff topology of sets), and also all the vertices of VG_θ vary continuously with θ near θ_i . It follows that the combinatorial changes that VG_θ undergoes at θ_i are just changes of the labels of all vertices of VG_θ that lie on $e_1 \cup e_2$ before crossing θ_i (and become points on $e_3 \cup e_4$ after crossing θ_i), and of their incident edges. Since VG_θ contains at most $O(kn)$ such vertices and edges, and since they can all be accessed in total $O(kn \log kn)$ time using our representation of this graph, we can update VG_θ (and EG) during processing of the $O(kn)$ orientations of type (iii)(a) in overall time $O(k^2 n^2 \log kn)$.

Finally, if θ_i is a critical orientation of type (iii)(b), let S be the side of B that becomes horizontal at θ_i , and let S_L, S_R be its endpoints. As in the above discussion, to update VG_θ at θ_i , one has to remove from this graph the auxiliary corners corresponding to all previous vertically maximal corners induced by contact pairs of the form (W, S_L) , and also remove all their incident edges, and then add new auxiliary corners corresponding to new vertically maximal corners induced by pairs of the form (W, S_R) , and also add their incident edges. Since these new auxiliary corners are not yet present in VG_θ , we need to create them and locate them on the boundary of K_θ in order to determine their adjacent vertices along this boundary. This can be accomplished by a standard and straightforward line-sweeping technique, running in time $O(kn \log kn)$. Since the number of type (iii)(b) orientations is only $O(k)$, processing of these orientations at Step (3) of the algorithm requires only $O(k^2 n \log kn)$ time.

In the final step (4), all we need to do is to pair each vertex $\xi = (u, (\Omega, \theta_i))$ of EG whose initial lifespan orientation is undefined, with a corresponding vertex $\xi' = (u, (\theta_1, \Omega))$ having the same label u and an undefined terminal lifespan orientation. For each such pair we form a new vertex $\xi^* = (u, (\theta_1, \theta_2))$ of EG and merge the two lists of edges incident to ξ, ξ' to a single list for ξ^* . Since there are only $O(kn)$ such pairs of vertices, this merging can be trivially accomplished in $O(kn)$ time. Hence the time required by the algorithm is at most $O(kn \lambda_6(kn) \log kn)$, as asserted. This completes the proof of the theorem. \square

Remarks. (1) As noted earlier, when $m \ll n$ (where m is the number of convex polygonal regions into which the polygonal obstacles are partitioned) the complexity of VG_θ for any fixed θ is only $O(km + n)$ which can be significantly smaller than $O(kn)$. We can exploit this fact to obtain a better bound on the complexity of our algorithm as follows.

For each of the $O(kn)$ contact pairs $O = (W, S)$ let J_O denote the angular interval of orientations θ of B at which an obstacle contact induced by O is possible (in the sense of Definition 2.1(c)). The $O(kn)$ endpoints of these intervals partition the angular space of orientations into $O(kn)$ disjoint intervals I_1, I_2, \dots, I_l . It follows from the arguments of [GRS] that within each of these

intervals I_j the set of possible obstacle contacts (in the sense of Definition 2.1(c)) remains constant, and that there are at most $km + n$ such contacts. Hence the number of pairs of contact pairs that can induced a double contact of B with an orientation in I_j is at most $\binom{km+n}{2} = O((km+n)^2)$. Moreover, as we cross from I_i to the next interval I_{j+1} , only one contact pair can induce a new obstacle contact, thus the number of new possible pairs of contacts within I_{j+1} is only $km+n$, and, repeating this argument over all intervals I_i , we obtain that the total number of double contacts of B is at most $\binom{km+n}{2} + (km+n)O(kn) = O((km+n)kn)$.

For each contact pair O , let t_O denote the number of bounding functions F_{OO} appearing in the envelopes $\Psi_{L;O}$, $\Psi_{R;O}$. Then the combinatorial complexity of $\Psi_{L;O}$, $\Psi_{R;O}$ is $O(\lambda_6(t_O))$, so that the total number of type (i) critical orientations is $\sum_O \lambda_6(t_O)$ which, as can be easily verified (using the bound in [Sh1]), is $O((km+n)\lambda_6(kn))$.

Similar arguments apply to all other types of critical orientations, and therefore imply that the complexity of our algorithm is at most $O((km+n)\lambda_6(kn) \log kn)$.

(2) The example given in [LS3] can be used to show that the combinatorial complexity of FP can be $\Omega(k^2n^2)$ in the worst case, so that our algorithm is close to being optimal among algorithms that calculate the entire space FP . See also a related result by O'Rourke [OR] for the case of a moving line segment.

3.5. Actual Motion Planning

Once EG is available, actual motion planning between any two given placements of B can be easily accomplished, as follows. Let $Z = (X, \theta)$, $Z' = (X', \theta')$ be two given free placements of B . We first calculate the nodes $\Phi(Z)$, $\Phi(Z')$ of EG (where Φ is the map defined in the paragraph following the proof of Proposition 2.1). It is easy to calculate these nodes in $O(kn)$ time in a straightforward manner following the definition of Φ . Next, determine by simple graph searching whether $\Phi(Z)$ and $\Phi(Z')$ belong to the same connected component of EG (this can be accomplished in time $O(kn\lambda_6(kn))$). If not, then by Theorem 2.2 no collision-free motion of B between placements Z and Z' is possible. Otherwise, let $\pi = (\xi_1 = \Phi(Z), \xi_2, \dots, \xi_t = \Phi(Z'))$ be a path in EG connecting $\Phi(Z)$ to $\Phi(Z')$. We transform π into a continuous semifree motion of B from Z to Z' as follows. First B is translated from Z to a placement on the edge ξ_1 of FP , as in the definition of the map Φ . Then, for each $i = 1, \dots, t-1$ we choose some orientation θ_i in the lifespan of the edge (ξ_i, ξ_{i+1}) of EG , and move B from its current position on ξ_i to a position on ξ_{i+1} in two substeps:

- (a) First B is moved along the edge ξ_i until its orientation becomes equal to θ_i . The exact nature of this motion depends on the type of corner u of K_θ that labels ξ_i . For example, if u is a convex corner of K_θ , induced by a contact pair (W, S) of type III, then the required motion of B is simply

a rotation of B about its corner S which remains touching the wall corner W . If u is a nonconvex corner of K_θ , induced by two contact pairs (W_1, S_1) , (W_2, S_2) of type I or II, B has to *glide* as it maintains the two contacts of S_1 against W_1 and of S_2 against W_2 . The resulting motion, known as a *glissette*, can be parametrized by a quartic polynomial equation (see [SS]). Finally, if u is an auxiliary corner v^* corresponding to some convex corner v of K_θ , let (W, S) be the contact pair inducing the edge of K_θ containing u , and let (W', S') be the type III pair inducing v . Then B has to move so that it maintains contact of S against W while its corner S' remains covertical with (and above) the wall corner W' . This motion is clearly also a fourth-degree glissette (where we can think of S' as maintaining contact with the vertical ray emanating upward from W').

- (b) After B has reached orientation θ_i , we next move it to the FP -edge ξ_{i+1} . By definition, VG_{θ_i} contains an edge (u, v) where u (resp. v) is the label of the corner of K_θ that induces ξ_i (resp. ξ_{i+1}). Thus the required motion of B can be accomplished by translating it along the straight boundary segment of K_θ , connecting u and v .

Finally, after reaching the last edge ξ_i , we complete the motion of B by first moving it along ξ_i until it reaches its final orientation θ' (as in (a) above), and then translate it back to Z' by applying the reverse of the canonical translation used in the definition of the map Φ .

Remark. The motion of B described above is semifree rather than free. This is a consequence of our choice of representation of FP , and resembles the motion obtained for a moving line segment [SiS].

References

- [At] M. Atallah, Dynamic computational geometry, *Proc. 24th Symp. on Foundations of Computer Science*, 1983, pp. 92-99.
- [BZ] B. K. Bhattacharya and J. Zorbas, Solving the two-dimensional findpath problem using a line-triangle representation of the robot. Tech. Rept., School of Computing Science, Simon Fraser University, Burnaby, B.C. V5A 1S6.
- [GRS] L. Guibas, L. Ramshaw, and J. Stolfi, A kinetic framework for computational geometry, *Proc. 24th IEEE Symp. on Foundations of Computer Science*, 1983, pp. 100-111.
- [HS] S. Hart and M. Sharir, Nonlinearity of Davenport-Schinzel sequences and of generalized path compression schemes, *Combinatorica* 6(2) (1986), 151-177.
- [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 Comput. Geom.* 1 (1986), 59-71.
- [KS] K. Kedem and M. Sharir, An efficient algorithm for planning collision-free translational motion of a convex polygonal object in 2-dimensional space amidst polygonal obstacles, *Proc. ACM Symp. on Computational Geometry*, 1985, pp. 75-80.
- [LS1] D. Leven and M. Sharir, An efficient and simple motion-planning algorithm for a ladder moving in two-dimensional space amidst polygonal barriers, *J. Algorithms* 8 (1987), 192-215.
- [LS2] D. Leven and M. Sharir, Planning a purely translational motion for a convex object in two-dimensional space using generalized Voronoi diagrams, *Discrete Comput. Geom.* 2 (1987), 9-31.

- [LS3] D. Leven and M. Sharir, On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space, *Discrete Comput. Geom.* **2** (1987) 255-270.
- [OSY1] C. Ó'Dúnlaing, M. Sharir, and C. Yap, Generalized Voronoi diagrams for a ladder: I. Topological analysis, *Comm. Pure Appl. Math.* **39** (1986), 423-483.
- [OSY2] C. Ó'Dúnlaing, M. Sharir, and C. Yap, Generalized Voronoi diagrams for a ladder: II. Efficient construction of the diagram, *Algorithmica* **2** (1987), 27-59.
- [OY] C. Ó'Dúnlaing and C. Yap, A "retraction" method for planning the motion of a disc, *J. Algorithms* **6** (1985), 104-111.
- [OR] J. O'Rourke, Lower bounds for moving a ladder, Tech. Rept. 85/20, Department of EECS, The Johns Hopkins University, 1985.
- [OWW] Th. Ottmann, P. Widmayer, and D. Wood, A fast algorithm for boolean mask operations, Rept. No. 112, Inst. f. Angewandte Mathematik und Formale Beschreibungsverfahren, D-7500 Karlsruhe, 1982.
- [SS] 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 Appl. Math.* **36** (1983), 345-398.
- [Sh1] M. Sharir, Almost linear upper bounds on the length of general Davenport-Schinzel sequences, *Combinatorica* **7** (1987), 131-143.
- [Sh2] M. Sharir, Improved lower bounds on the length of Davenport-Schinzel sequences, *Combinatorica* **8** (1988), to appear.
- [SiS] S. Sifrony and M. Sharir, A new efficient motion-planning algorithm for a rod in two-dimensional polygonal space, *Algorithmica* **2** (1987), 367-402.
- [Sz] E. Szemerédi, On a problem by Davenport and Schinzel, *Acta Arith.* **25** (1974), 213-224.

Received October 25, 1986, and in revised form May 25, 1987.