Simplified Voronoi Diagrams*

John Canny
Bruce Donald

87-879
November 1987

Department of Computer Science
Cornell University
Ithaca, New York 14853-7501

*This report appeared in the *Proceedings of the 3rd ACM Symposium on Computational Geometry*, Waterloo, Ontario, (June 1987). It will also appear in the journal *Discrete and Computational Geometry.*

Simplified Voronoi Diagrams

John Canny

Bruce Donald

Abstract:

We are interested in Voronoi diagrams as a tool in robot path planning, where the search for a path in an r dimensional space may be simplified to a search on an $r-1$ dimensional Voronoi diagram. We define a Voronoi diagram V based on a measure of distance which is not a true metric. This formulation has lower algebraic complexity than the usual definition, which is a considerable advantage in motion planning problems with many degrees of freedom. In its simplest form, the measure of distance between a point and a polytope is the maximum of the distances of the point from the half-spaces which pass through faces of the polytope. More generally, the measure is defined in configuration spaces which represent rotation. The Voronoi diagram defined using this distance measure is no longer a strong deformation retract of free space, but it has the following useful property: any path through free space which starts and ends on the diagram can be continuously deformed so that it lies entirely on the diagram. Thus it is still complete for motion planning, but it has lower algebraic complexity than a diagram based on the euclidean metric.

Acknowledgements. This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's Artificial Intelligence research is provided in part by the Office of Naval Research under Office of Naval Research contract N00014–81–K–0494 and in part by the Advanced Research Projects Agency under Office of Naval Research contracts N00014–85–K–0124 and N00014–82–K–0334. John Canny was supported by an IBM fellowship. Bruce Donald was funded in part by a NASA fellowship administered by the Jet Propulsion Laboratory.

Authors' current addresses:
John Canny, Computer Science Division, University of California, Berkeley, CA 94720.
Bruce Donald, Computer Science Department, Upson Hall, Cornell University, Ithaca, NY 14853.

Proc. 3rd ACM Symposium on Computational Geometry
Waterloo, Ontario (June 1987)

1 Introduction

The Voronoi diagram has proved to be a useful tool in a variety of contexts in computational geometry. Our interest here is in using the diagram to simplify the planning of collision-free paths for a robot among obstacles, the so-called generalized movers' problem. The Voronoi diagram, as usually defined, is a *strong deformation retract* of free space so that free space can be continuously deformed onto the diagram. This means that the diagram is complete for path planning, i.e. Searching the original space for paths can be reduced to a search on the diagram. Reducing the dimension of the set to be searched usually reduces the time complexity of the search. Secondly, the diagram leads to robust paths, i.e. paths that are maximally clear of obstacles.

The Voronoi diagram generated by a set of points in a Euclidean space partitions the space into convex regions which have a single nearest point under some (usually L_2) metric. A generalized Voronoi diagram can be defined for points and line segments in the plane (Lee and Drysdale, 1981) which partitions the plane into (generally non-convex) regions. In both cases the diagram is defined to be the set of points equidistant from two or more generators under the appropriate metric. This construction has proved to be useful for motion planning among a set of obstacles in configuration space (see Ó'Dúnlaing and Yap (1982), Ó'Dúnlaing, Sharir, and Yap (1984), Yap (1984), and the textbook of Schwartz and Yap (1986) for an introduction and review of the use of Voronoi diagrams in motion planning). Its virtue for motion planning is that the diagram is a strong deformation retract of free space, i.e. the space outside the obstacles can be continuously deformed onto the diagram. To find a path between two points in free space, it suffices to find a path for each point onto the diagram, and to join these points with a path that lies wholly on the diagram.

The simplified diagram has lower algebraic complexity than the L_2 diagram. For example, in $\Re^3$, the L_2 diagram about polyhedral obstacles consists of quadric sheets; the simplified diagram is piecewise linear. In $\Re^2$, the simplified diagram for polygonal obstacles is a graph of straight lines, see figure 5. In general, the simplified diagram has the same degree as the algebraic obstacle constraints. However, it may not have linear size in the worst case.

One useful aspect of the simplified Voronoi diagram is that it is naturally defined for the six dimensional configuration space of an arbitrary 3D polyhedron moving amidst 3D polyhedral obstacles. Our definition elaborates a suggestion of Donald (1984) and Canny (1985), who describe certain Voronoi-like properties of the algebraic set $\bigcup_{i \neq i'} \ker(C_i - C_{i'})$ for a set of algebraic constraints $\{ C_i \}$. In this paper, we consider the configuration space of a polyhedral object with translational and rotational degrees of freedom. The simplified Voronoi diagram has the same algebraic complexity as the resulting configuration space obstacle boundaries. The completeness property holds for the simplified diagram when the defining algebraic obstacle constraints in the configuration space have unit gradients. The diagram has the same degree as these normalized constraints. (The *degree* of the diagram is the degree of the defining equations). Thus in $\Re^2$ and $\Re^3$ the simplified diagram has

degree one whereas the Euclidean diagram has degree 2. However, note that the Euclidean diagram in $\Re^3$ has curves of degree 4 and vertices of degree eight, whereas the simplified diagram is piecewise linear. In the configuration space $\Re^2 \times S^1$ of a planar polygon, the simplified diagram has degree 3, whereas the Euclidean diagram has degree 6. In the six-dimensional configuration space $\Re^3 \times SO(3)$ of a 3D polyhedron, both the simplified and Euclidean diagram have degree 10. A one-dimensional skeleton of the simplified diagram (in 6D) can be computed in time $O(n^7 \log n)$; this is the cost of computing an obstacle avoiding path along the diagram.

The completeness proof in this paper holds for the normalized six-dimensional case, and specializes to the lower-dimensional cases. The simplified diagram can also be defined for unnormalized configuration space constraints. In the lower-dimensional configuration spaces, the constraints may be normalized with no increase in algebraic complexity. In the six-dimensional case, a quaternion representation of rotation gives constraints of degree three in the configuration parameters. In fact the constraints are simultaneously quadric in the quaternion components and linear in the position components. Some of these constraints do not a priori satisfy the normalization condition, but they can be normalized by dividing by a polynomial factor. Since this increases the degree of the equations defining the diagram (to an effective degree of 10, the same as the Euclidean diagram), we suggest instead that they be left unnormalized, and we offer the completeness proof in this paper as heuristic evidence that the property holds for a reasonable class of unnormalized systems of constraints.

2 Object-Obstacle Constraints

We briefly derive conditions for overlap of two polyhedral objects A and B. A more complete derivation of an equivalent condition is given in Canny (86). The form we derive here is different from (Canny 86) in that it uses a local test for non-overlap, rather than overlap. We assume first that A and B are convex, and then generalize to the non-convex case by taking the conjunction of pairwise non-overlap predicates between convex pieces.

The overlap predicates in Canny (86) generate a shallow (depth 2) AND-OR predicate tree, whose root is a disjunction. It will be advantageous to make the predicate tree as deep as possible and it is also desirable for the root to be a conjunction. So instead we use the following test based on conditions for non-overlap.

Definition.

For any face f of a convex polyhedron A, the *affine hull* $\overline{f}$ of f is the plane which contains f.

The affine hull of a face f defines two closed half-spaces, one of which contains A. We call this half-space the *interior half-space*, and denote it $\overline{f}^-$.

Finally, we define the *wedge* of an edge e of A as the intersection of the two interior half-spaces of the faces which cobound e. The wedge of e is denoted $\hat{e}$, and it contains A.

Lemma 2.1 *Two convex polyhedra A and B are non-overlapping iff either all edges of A are outside some wedge of an edge of B, or all edges of B are outside of some wedge an edge of A.*

Proof. This condition is clearly sufficient for non-overlap, by convexity of A and B. Conversely, if A and B are disjoint, then there is a (not necessarily unique) non-zero shortest vector between them. Let p_A and p_B be the end-points of this vector in A and B respectively. If one of these points lies in the interior of a face f, then the test succeeds for any wedge of an edge in the boundary of f. If one of the points lies in the interior of an edge e, then the test succeeds for $\hat{e}$

The only case remaining is where p_A and p_B are both vertices. Let f be a face adjacent to p_A, and such that p_B lies outside the interior half-space $\overline{f}^-$, (there must be at least one such f, or p_B would be contained in A). Then if B is also outside this half-space, the test succeeds for the wedge of any edge that cobounds f.

Otherwise, let W_{p_B} be the intersection of the wedges of all edges that cobound p_B. Let S be the plane passing through p_B and normal to the vector $(p_B - p_A)$. S defines two closed half-spaces, one (S_A) containing A and the other (S_B) containing both B and W_{p_B}, i.e. S separates A and B. Then the intersection $(W_{p_B} \cap \overline{f})$ lies in $(S_B \cap \overline{f})$. Let p_0 be the closest point in $(W_{p_B} \cap \overline{f})$ to $(S \cap \overline{f})$. Then p_0 is in the boundary of some wedge $\hat{e}$ of an edge e that cobounds p_B. Now $(\hat{e} \cap \overline{f}) \subset (S_B \cap \overline{f})$ and so by projection from p_B, $(\hat{e} \cap \overline{f}^-) \subset (S_B \cap \overline{f}^-)$. But $A \subset (S_A \cap \overline{f}^-)$, so $(\hat{e} \cap A) = \emptyset$, and the test succeeds for $\hat{e}$. □

Thus we can define the following predicate for non-overlap $F_{A,B}$ of A and B from the above test:

$$F_{A,B} = \Big(\bigvee_{\substack{e_j \in \\ \text{edges}(B)}} \bigwedge_{\substack{e_i \in \\ \text{edges}(A)}} F_{e_i,\hat{e}_j} \Big) \bigvee \Big(\bigvee_{\substack{e_i \in \\ \text{edges}(A)}} \bigwedge_{\substack{e_j \in \\ \text{edges}(B)}} F_{\hat{e}_i,e_j} \Big) \tag{1}$$

where $F_{\hat{e}_i,e_j} = ((\hat{e}_i \cap e_j) = \emptyset)$. The corresponding condition for overlap is in the form of a conjunction of disjunctions, as desired, and this is the form we will use in our development. If the object consists of several convex pieces A_i, as do the obstacles B_j, then the non-overlap predicate is the conjunction of pairwise predicates

$$F = \bigwedge_i \bigwedge_j F_{A_i,B_j} \tag{2}$$

We must now decompose the non-overlap predicate for a wedge $\hat{e}_A$ of A and a edge e_B of B into simple geometric predicates that can be computed directly. These geometric predicates are $A_{f,p}$, $C^+_{e_A,e_B}$ and $C^-_{e_A,e_B}$. $A_{f,p}$ indicates non-overlap of a vertex p of B and the interior half-space of a face f of A and is given by

$$A_{f,p} = (\mathbf{n}_f \cdot p - c_f > 0) \tag{3}$$

where $\mathbf{n}_f$ is the outward normal of f, and c_f is its distance from the origin. For $C^+_{e_A,e_B}$ and $C^-_{e_A,e_B}$ we need the following definitions: Let d_A and d_B be the vector directions of e_A and e_B respectively, and let p_A be any point on e_A. Let and H and T be the head and tail vertex respectively of e_B, then we have

$$C^+_{e_A,e_B} = (H - p_A) \cdot (d_A \times d_B) > 0 \tag{4}$$

$$C^-_{e_A,e_B} = (H - p_A) \cdot (d_A \times d_B) < 0$$

Now we can define $F_{\hat{e}_A,e_B}$ in terms of the above predicates, and L and R which are the left and right faces respectively, which cobound e_B, (left and right are determined here by viewing e_B from outside $\hat{e}_B$ with d_B upward).

Lemma 2.2 *The following predicate indicates non-overlap of the wedge $\hat{e}_A$ and the edge e_B.*

$$\begin{aligned} F_{\hat{e}_A,e_B} \;=\; & ((A^+_{L,H} \wedge (A^+_{L,T} \vee (A^+_{R,T} \wedge C^+_{e_A,e_B}))) \\ & \vee (A^+_{R,H} \wedge (A^+_{R,T} \vee (A^+_{L,T} \wedge C^-_{e_A,e_B})))) \end{aligned} \tag{5}$$

Proof The proof is by analysis of all possible predicate values. Firstly, the predicate $A^+_{L,H}$ requires the head of e_B to be above the plane of the left face of e_A. For the rest of this paragraph, we assume that $A^+_{L,H}$ is true, so that H is above the plane of L. The subcases are itemized below:

- If $A^+_{L,T}$ is true, then the tail T of e_B is also above the plane of L, i.e. the entire edge e_B is outside of $\hat{e}_A$ and the predicate correctly returns true.
- If $A^+_{L,T}$ is false, and so is $A^+_{R,T}$, then the vertex T lies inside $\hat{e}_A$, and the predicate $F_{\hat{e}_A,e_B}$ correctly returns false.
- The only case remaining is where $A^+_{L,T}$ is false and $A^+_{R,T}$ is true, so that both vertices are outside of $\hat{e}_A$, and H is above the plane of L and T is below that plane. This case is illustrated in figure 1. In this case $d_A \times d_B$ is a vector which points out of $\hat{e}_A$, and is normal to d_A and d_B. Then e_A is outside of $\hat{e}_A$ if and only if the inner product of $(H - p_A)$ with this vector is positive, a condition which is indicated by $C^+_{e_A,e_B}$.

The subcases for $A^+_{R,H}$ true are similar, with left and right faces interchanged. The only case remaining is where both $A^+_{L,H}$ and $A^+_{R,H}$ are false, but here the point H is inside the object and the predicate $F_{\hat{e}_A,e_B}$ correctly returns a false value. □

The predicate $F_{\hat{e}_A,e_B}$ can be written in the equivalent form:

$$\begin{aligned} F_{\hat{e}_A,e_B} \;=\; & (A_{L,H} \vee A_{R,H}) \wedge (A_{L,T} \vee A_{R,T}) \\ & \wedge (A_{L,H} \vee A_{R,T} \vee C^+_{e_A,e_B}) \wedge (A_{R,H} \vee A_{L,T} \vee C^-_{e_A,e_B}) \end{aligned} \tag{6}$$

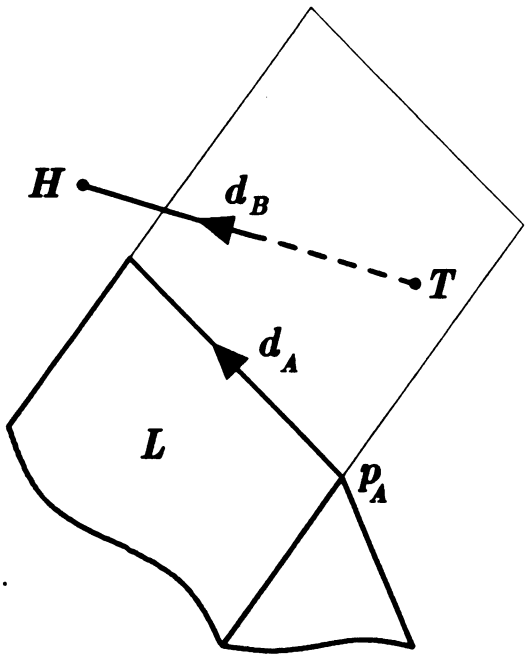

Figure 1: Position of edge e_B outside the wedge $\hat{e}_A$, where inside/outside is determined by $C^+_{e_A,e_B}$

There is a similar form for the predicate $F_{e_A,\hat{e}_B}$. From (1), (2) and (6) we see that the overall form of the non-overlap predicate can be written as

$$F = \bigwedge_i \bigvee_j \bigwedge_k \bigvee_l (C_{ijkl} > 0) \tag{7}$$

3 The Voronoi Diagram

For motion planning, the configuration of object A is variable. For a polyhedron in three dimensional space, the configuration contains a position $\mathbf{x} \in \Re^3$ and an orientation $\mathbf{q} \in SO(3)$ component, where $SO(3)$ is the group of three-dimensional rotations. Thus the face normals, vertex locations, and edge directions of A are all functions of $\mathbf{x}$ and $\mathbf{q}$. The predicate (7) is now also a function of the configuration $(\mathbf{x}, \mathbf{q})$, i.e.

$$F(\mathbf{x}, \mathbf{q}) = \bigwedge_i \bigvee_j \bigwedge_k \bigvee_l (C_{ijkl}(\mathbf{x}, \mathbf{q}) > 0) \tag{8}$$

The forms of the functions $C_{ijkl}(\mathbf{x}, \mathbf{q})$ are given explicitly in Canny (1986), and they are algebraic if, say, a quaternion representation of rotation is used. The set of overlap configurations is called the *configuration obstacle* and is denoted $CO = \{(\mathbf{x}, \mathbf{q}) | \neg F(\mathbf{x}, \mathbf{q})\}$. It may be thought of as a physical obstacle in configuration space to be avoided by a path

planner. We now observe that by letting positive real values represent logical one, and non-positive values represent logical zero, that the min function implements logical AND, and the max function implements logical OR. Thus an equivalent form to (8) is

$$F(\mathbf{x},\mathbf{q}) = \left(\left(\min_i (\max_j (\min_k (\max_l C_{ijkl}(\mathbf{x},\mathbf{q})))) \right) > 0 \right) \tag{9}$$

which suggests that the quantity

$$\rho(\mathbf{x},\mathbf{q}) = \min_i (\max_j (\min_k (\max_l C_{ijkl}(\mathbf{x},\mathbf{q})))) \tag{10}$$

can be used as a measure of distance from the configuration obstacle, because it varies continuously through configuration space, is positive at configurations outside CO, and non-positive at configurations inside CO. Thus the configuration obstacle can be rewritten as $CO = \{\mathbf{p} \mid \rho(\mathbf{p}) \leq 0\}$, and its complement, the set of points in free space can be written $F = \{\mathbf{p} \mid \rho(\mathbf{p}) > 0\}$.

In order to define the Voronoi diagram under the distance measure ρ, we need a notion of closest feature. The closest features to a configuration $(\mathbf{x},\mathbf{q})$ are those C_{ijkl} which are *critical* in determining the value of $\rho(\mathbf{x},\mathbf{q})$, that is, small changes in the value of C_{ijkl} cause indentical changes in the value of ρ.

Definition.

A constraint $C_{i_0 j_0 k_0 l_0} \in \{C_{ijkl}\}$ is *critical* at a configuration $(\mathbf{x},\mathbf{q})$ if the value of $C_{i_0 j_0 k_0 l_0}(\mathbf{x},\mathbf{q})$ equals the maximum (or minimum) value of every max (resp. min) ancestor of $C_{i_0 j_0 k_0 l_0}$ in the min-max tree in (10). i.e.

$$C_{i_0 j_0 k_0 l_0}(\mathbf{x},\mathbf{q}) = \max_l C_{i_0 j_0 k_0 l}(\mathbf{x},\mathbf{q}) = \min_k (\max_l C_{i_0 j_0 k l}(\mathbf{x},\mathbf{q})) = \ldots \tag{11}$$

Now we have

Definition.

The *simplified Voronoi diagram* V is the set of configurations in free space F at which at least two distinct constraints are critical.

It should be clear from the definition of criticality that V is semi-algebraic if the constraints C_{ijkl} are algebraic. V is closed as a subset of free space, although it is not closed in configuration space. Notice that V has no interior, since it is contained inside a finite set of bisectors, each of which has no interior. A bisector is the zero set of $(C_{i_0 j_0 k_0 l_0} - C_{i_1 j_1 k_1 l_1})$ for some pair of distinct constraints. It will prove useful to subdivide the Voronoi diagram into two parts:

Definition.

The *concave part* of the Voronoi diagram V denoted conc(V) is the set of configurations in F where two distinct constraints are critical, and the lowest common ancestor of these constraints in the min-max tree of (10) is a min-node.

The *convex part* of the Voronoi diagram V denoted conv(V) is the set of configurations in F where two distinct constraints are critical, and the lowest common ancestor of these constraints is a max-node.

Notice that these two definitions are not mutually exclusive, because there may be points where more than two constraints are critical, and which satisfy both definitions. Thus conv(V) and conc(V) may overlap.

4 Completeness for Motion Planning

Our key result is that any path in F with endpoints in V can be deformed (in F) to a path with the same endpoints lying entirely in V. We start with a path $p : I \rightarrow \Re^3 \times SO(3)$ lying in free space, $p(I) \subset F$, where $I = [0,1]$ is the unit interval.

First we assume *wlog* that p intersects V at a finite number of points. We can do this because, as defined in the previous section, V is a semi-algebraic set, and by Whitney's (1957) result, it can be split into a finite number of manifolds, or *strata*. Since V has no interior, all these manifolds have codimension at least one. For *any* path p there is a path p' arbitrarily close to p, which is an embedding of I, and so $p'(I)$ is a 1-manifold. Almost every perturbation of p' intersects all of the strata transversally, and therefore at a finite number of points. We can choose such a perturbation to be arbitrarily small, in particular, smaller than the minimum distance from $p'(I)$ to CO. Such a perturbation gives a new path p'' which is path homotopic to p, and which has finite intersection with V.

So we assume that p has m intersections with V, and that these occur at points $p(x_i)$ with $x_1, \ldots x_m \in I$, and $x_1 = 0$ and $x_m = 1$. We then break each interval $[x_i, x_{i+1}]$ in half, giving us two intervals sharing an endpoint. Thus we now have $2m-2$ intervals each of which intersects V at only one endpoint. Below we give homotopies for each of these intervals which continuously deform the image of the interval onto V. Since all these homotopies agree at their endpoints, they can be pasted together to give us a global homotopy which deforms p onto V. For simplicity, we assume the path segment is parametrized in the range $I = [0,1]$ and that $p(0) \in V$.

The motion constraints C_{ijkl} are either A or C predicates, (3) and (4), and all can be written in the general form below, called the parametrized plane equations by Lozano-Perez (1983).

$$C_{ijkl}(\mathbf{x}, \mathbf{q}) \;=\; \mathbf{N}_{ijkl}(\mathbf{q}) \cdot \mathbf{x} + c_{ijkl}(\mathbf{q}) \tag{12}$$

Where $\mathbf{N}_{ijkl}(\mathbf{q}) \in \Re^3$ and $c_{ijkl}(\mathbf{q}) \in \Re$ are both continuous functions of $\mathbf{q}$. We assume that the C_{ijkl} are normalized so that $|\mathbf{N}_{ijkl}(\mathbf{q})| = 1$ for all $\mathbf{q}$. Our objective is to continuously deform the path p onto the diagram, and we use the $\mathbf{N}_{ijkl}$ as 'normals' to push a point on

the path $p(I)$ away from the critical C_{ijkl}. We assume that the set of positions is bounded by some set of constraint 'walls' of the same form as (1), so that a point can be displaced only a finite distance in free space. We also assume that the workspace has unit diameter.

General position assumptions. The construction requires the following general position assumptions. The second assumption requires an arbitrarily small perturbation of the constraints that are used to define the diagram, and this should be done as a preprocessing step before computation of the diagram. First, suppose C_{ijkl} is type C predicate (4). Then $\mathbf{N}_{ijkl}(\mathbf{q}) = d_A(\mathbf{q}) \times d_B$. To normalize $\mathbf{N}_{ijkl}$, we must divide by its magnitude, which must remain non-zero. Hence the set of configurations

$$\{\, \mathbf{q} \mid d_A(\mathbf{q}) \times d_B = 0 \,\} \subset SO(3) \tag{13}$$

must not intersect the image of p. However, (13) is clearly of codimension 2 in $SO(3)$, and hence by Sard's lemma there is always an arbitrarily small perturbation of p which avoids (13).

Similarly, the set

$$\{\, (\mathbf{x}, \mathbf{q}) \mid (\mathbf{N}_{ijkl}(\mathbf{q}) = \pm \mathbf{N}_{i'j'k'l'}(\mathbf{q})) \,\} \tag{14}$$

is also of codimension 2 in $\Re^3 \times SO(3)$, and the proof is given below for the "+" case. For the $-$ case, the below argument can be applied by simply negating one of the normals.

The singular set where two (unit) normals are equal has codimension two for the following reason. We consider a map from $\Re^3 \times SO(3)$ to the product $S^2 \times S^2$ which represents the values of the two normals. The diagonal of this space is the set where the two normals agree, and has codimension 2. If the maps are generic (specifically, if they are transversal to the diagonal set), then the preimage of the diagonal set (which is the "bad" set) will also have codimension 2.

We must define two different homotopies depending on whether $p(0) \in \mathrm{conc}(V)$, which is the simplest case, or $p(0) \in \mathrm{conv}(V)$.

Notice that since $p(0)$ is the only point on the path which is in V, there is exactly one constraint which is critical at all configurations in $p(0, 1]$ (since constraints change value continuously along the path and for another constraint to become critical, it must first equal the first constraint, which can only occur at points on the diagram). Let the critical constraint be C_{ijkl}. We define a homotopy $J_0 : I \times I \to \Re^3 \times SO(3)$ as

$$J_0(t, u) = p(t) + u\mathbf{N}_{ijkl}(\pi_q(p(t))) \tag{15}$$

where $\pi_q(p(t))$ is the orientation component of $p(t)$ and the addition symbol means we add the vector quantity $u\mathbf{N}_{ijkl}(\pi_q(p(t))$ to the position component of $p(t)$. The deformation above pushes points beyond the diagram, so we define a second homotopy $J_1 : I \times I \to F$.

$$J_1(t, u) = J_0(t, \min(u, u_c(t))) \tag{16}$$

$$\text{where} \quad u_c(t) = \inf\{\, s \mid \; J_0(t,s) \in \text{conc}(V)\,\}$$

Note that u_c is bounded above by 1 since the workspace has unit diameter. Recall from the definition of homotopy, that J_1 is a homotopy of p and a path p' if J_1 is continuous and $J_1(t,0) = p(t)$ and $J_1(t,1) = p'(t)$. The homotopy J_1 suffices to map paths with one endpoint in conc(V) onto V:

Lemma 4.1 *Let $p : I \rightarrow \Re^3 \times SO(3)$ be a path having $p(0) \in V$ and no other points in V. Then J_1 is a homotopy of p and a path p' such that $p'(I) \subset \text{conc}(V)$. Furthermore $p'(0) = p(0)$ if $p(0) \in \text{conc}(V)$.*

Proof. From the definition of J_1 we have $J_1(t,0) = p(t)$ and $J_1(t,1) \in \text{conc}(V)$. Also if $p(0) \in \text{conc}(V)$ then $u_c(0) = 0$ so $p'(0) = p(0)$. It remains to show that J_1 is continuous. First we notice that J_0 is continuous. Continuity of J_1 follows if we can show that $u_c(t)$ is continuous. Now $J_0(t, u_c(t))$ is contained in the zero sets of all bisectors $\{(C_{i'j'k'l'} - C_{ijkl})\}$. Let $u'_c(t)$ be a deformation onto a *particular* bisector:

$$u'_c(t) = \inf\{u \mid \; J_0(t,u) \in \ker(C_{i'j'k'l'} - C_{ijkl})\} \tag{17}$$

then u'_c is continuous because by definition

$$(C_{ijkl} - C_{i'j'k'l'})(J_0(t, u'_c(t))) = 0 \tag{18}$$

and differentiating with respect to t, we obtain

$$\tfrac{\partial}{\partial t}(C_{ijkl} - C_{i'j'k'l'})(J_0(t,u'_c)) + \left(\tfrac{\partial u'_c}{\partial t}\right) \tfrac{\partial}{\partial u'_c}(C_{ijkl} - C_{i'j'k'l'})(J_0(t,u'_c)) \;=\; 0 \tag{19}$$

which can be rearranged to yield

$$\frac{\partial u'_c}{\partial t} \;=\; -\frac{\tfrac{\partial}{\partial t}(C_{ijkl} - C_{i'j'k'l'})(J_0(t,u'_c))}{\tfrac{\partial}{\partial u'_c}(C_{ijkl} - C_{i'j'k'l'})(J_0(t,u'_c))} \tag{20}$$

and therefore $\frac{\partial}{\partial t}u'_c(t)$ is finite, because the denominator above is non-zero by our general position assumption. Now we observe that $u_c(t)$ can be constructed by pasting together segments of $u'_c(t)$ for various bisectors $\in \text{conc}(V)$, and we must show that they agree at their endpoints. The proof is by contradiction. Suppose we had

$$u_c(t) = \begin{cases} u'_c(t) & \text{for } t \in (t_0, t_1]; \\ u''_c(t) & \text{for } t \in (t_1, t_2); \end{cases} \tag{21}$$

$$\text{with} \quad u'_c(t_1) \neq u''_c(t_1)$$

for u' and u'' derived from distinct bisectors. Then since u_c is the minimum of all such u, we must have $u'_c(t_1) < u''_c(t_1)$

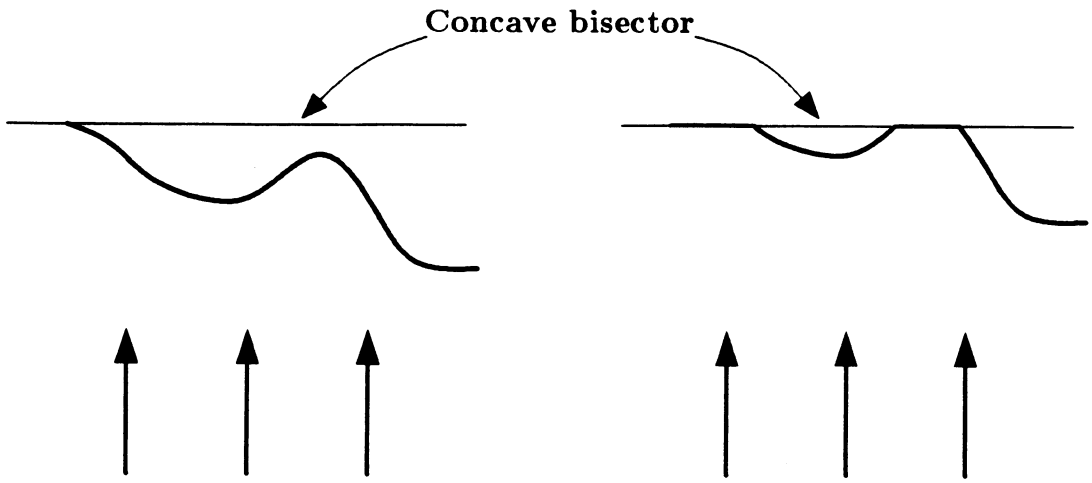

Figure 2: Illustration of the homotopy of lemma 4.2

But this means that as u increases, $J_0(t_1, u)$ crosses two bisectors in conc(V) between C_{ijkl} and other constraints. This is impossible because all constraints C have $|\mathbf{N}(\mathbf{q})| = 1$, and it follows that

$$\frac{\partial}{\partial u} C_{i'j'k'l'}(J_0(t_1, u)) \leq 1 \tag{22}$$

$$\frac{\partial}{\partial u} C_{ijkl}(J_0(t_1, u)) = 1$$

That is, all constraints increase no faster than C_{ijkl} with the deformation parameter u. By our general position assumption if $C_{ijkl} \neq C_{i'j'k'l'}$, then the inequality in (22) is strict, and so if the clause $C_{i'j'k'l'}$ becomes critical at u'_c, it shares a min node lowest ancestor with C_{ijkl}. Since all constraints in the tree increase more slowly than C_{ijkl}, the value of this min node will be less than the value of C_{ijkl} for $u > u'_c(t_1)$. Therefore C_{ijkl} cannot be critical for any $u > u'_c(t_1)$, contradicting the assumption that $u''_c(t_1)$ is distinct from $u'_c(t_1)$. So all bisectors in conc(V) between C_{ijkl} and other constraints agree at their endpoints, and $u_c(t)$ is continuous by pasting. This shows that J_1 is continuous. □

The homotopy described above is illustrated in figure 2. The direction of the gradient of the distance function is shown by the arrows. Each point on the path moves in the gradient direction until it hits a concave bisector.

If $p(0) \in \text{conv}(V)$ the situation is more complicated, because the deformation J_1 pushes $p(0, 1]$ away from $p(0)$. To correct this, we first compress the first half of p to a point:

Lemma 4.2 *Let $p : I \to F$ be a path in free space. Then p is homotopic to a path p' such that*

(i) $p'([0, \frac{1}{2}]) = p(0)$
(ii) $p'(1) = p(1)$
(iii) $p'(I) = p(I)$

Proof. The required homotopy is

$$J(t,u) = \begin{cases} p(0) & \text{if } t \leq \frac{1}{2}u; \\ p\left(\frac{2t-u}{2-u}\right) & \text{otherwise;} \end{cases} \tag{23}$$

so $J(t,0) = p(t)$, and $J(t,1) = p'(t)$. □

We apply the homotopy of lemma 4.2 to p to give us a path p'. Applying the homotopy J_1 to the path segment $p'|_{[\frac{1}{2},1]}$ continuously deforms this path segment onto conc(V). Then we define a new homotopy which slowly "unravels" $p'|_{[0,\frac{1}{2}]}$ from $p(0)$. All the points in this homotopy have the same orientation, and for each value of the deformation parameter u, the path consists of a finite number of straight line segments.

We now construct a third homotopy. The construction is inductive and we start with a homotopy that gives us two straight line segments. If the orientation at the configuration $p(0)$ is $\mathbf{q_0}$, every point on the joining path will also have orientation $\mathbf{q_0}$. We define two vectors in position space $\mathbf{N_0}$ and $\mathbf{M_0}$ which will be used to define the joining path segment.

Let C_{ijkl} and $C_{i'j'k'l'}$ be the two constraints that are critical at $p(0)$, then $\mathbf{N_0}$ lies in the plane of the bisector of C_{ijkl} and $C_{i'j'k'l'}$. $\mathbf{N_0}$ is normalized so that $\mathbf{N_0} \cdot \mathbf{N}_{ijkl}(\mathbf{q_0}) = 1$ and it follows that $\mathbf{N_0} \cdot \mathbf{N}_{i'j'k'l'}(\mathbf{q_0}) = 1$. A second vector $\mathbf{M_0}$ is chosen so that $\mathbf{M_0} + \mathbf{N}_{ijkl}(\mathbf{q_0}) = \mathbf{N_0}$, and so:

$$\mathbf{N_0} = \frac{\mathbf{N}_{ijkl} + \mathbf{N}_{i'j'k'l'}}{1 + \mathbf{N}_{ijkl} \cdot \mathbf{N}_{i'j'k'l'}} \tag{24}$$

$$\mathbf{M_0} = \frac{\mathbf{N}_{i'j'k'l'} - (\mathbf{N}_{ijkl} \cdot \mathbf{N}_{i'j'k'l'})\mathbf{N}_{ijkl}}{1 + \mathbf{N}_{ijkl} \cdot \mathbf{N}_{i'j'k'l'}} \tag{25}$$

where $\mathbf{N}_{ijkl}$ (shorthand for $\mathbf{N}_{ijkl}(\mathbf{q_0})$) is the normal vector to the critical constraint C_{ijkl} at orientation $\mathbf{q_0}$, and $\mathbf{N}_{i'j'k'l'} = \mathbf{N}_{i'j'k'l'}(\mathbf{q_0})$ is the vector normal to the critical constraint $C_{i'j'k'l'}$. (Note that $\mathbf{N}_{ijkl}$ cannot equal $-\mathbf{N}_{i'j'k'l'}$ by our general position assumptions).

We can now define the joining homotopy $J_2 : I \times I \to \Re^3 \times SO(3)$ which deforms the path segment $p'|_{[0;\frac{1}{2}]}$ (with p' reparametrized so that its domain is I):

$$J_2(t,u) = \begin{cases} p(0) + 2t\mathbf{N_0}, & \text{if } t \in [0, \frac{1}{2}u]; \\ p(0) + u\mathbf{N_0}, & \text{if } t \in [\frac{1}{2}u, 1 - \frac{1}{2}u]; \\ p(0) + (2-2t)\mathbf{M_0} + u\mathbf{N}_{ijkl}(\mathbf{q_0}), & \text{if } t \in [1 - \frac{1}{2}u, 1]; \end{cases} \tag{26}$$

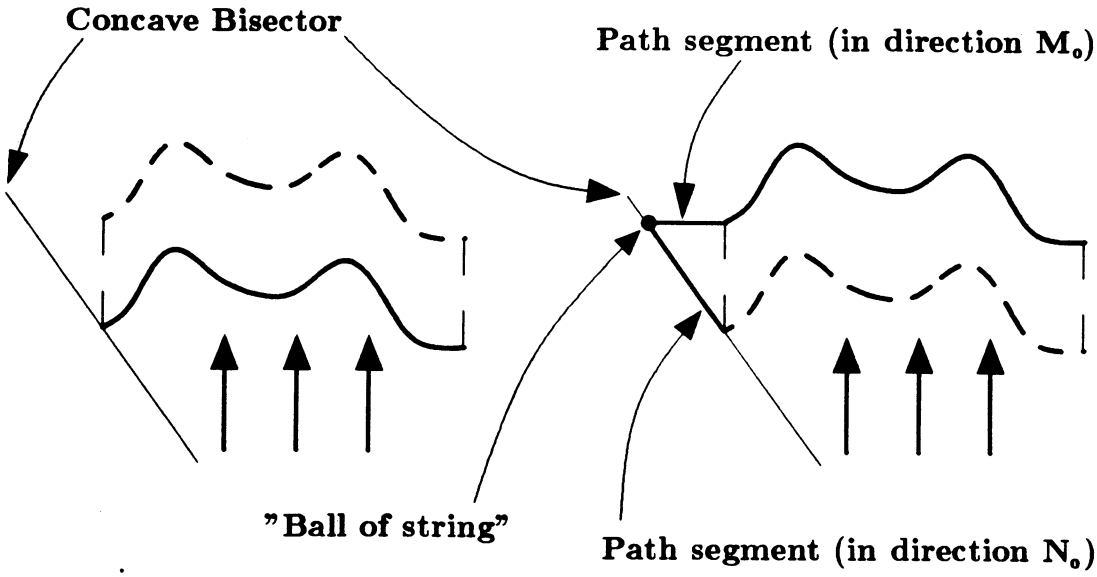

Figure 3: Homotopy to continuously link a deforming path to a point on a convex bisector

The action of this homotopy on a path segment is illustrated in figure 3. All points not on a bisector move in the direction of the gradient of the distance function. Points which are on the bisector do not move, except the single corner point. This point may be thought of as a "ball of string", which continuously unravels and allows points on the two straight line segments to move as described above.

Now $\mathbf{N_0}$ lies in the plane of the bisector of the two constraints that are critical at $p(0)$, but it is possible that as u increases, $J_2(\frac{1}{2}, u)$ leaves the convex part of the diagram before reaching the concave part. That is, there may be bends in the convex part of the diagram which must be tracked. We must therefore stop the deformation at this point by defining

$$u_v = \sup\{u \mid C_{ijkl} \text{ is critical throughout } J_2(\tfrac{1}{2}, [0, u]) \tag{27}$$

and once again we define a homotopy which stops points when they reach the diagram:

$$J_3(t, u) = J_2(t, \min(u, u_v, u_c(t))) \tag{28}$$

$$\text{where} \quad u_c(t) = \inf\{\, u \mid J_2(t, u) \in \text{conc}(V)\,\}.$$

and then we have:

Lemma 4.3 *Let $p : I \to \Re^3 \times SO(3)$ be a path having $p(I) = p(0) \in \text{conv}(V)$. Then J_3 is a homotopy of p and a path p' such that*

(i) $p'(0) = p(0)$

(ii) $p'([0, 1 - \frac{1}{2}u_v]) \subset \text{conv}(V)$

(iii) $p'((1 - \frac{1}{2}u_v, 1]) \cap V \subset \text{conc}(V)$

Proof.

Parts (i) and (ii) follow immediately from the definition of J_2 and J_3. Part (iii) states that points in the interval $(1-\frac{1}{2}u_v, 1]$ that are mapped into the diagram by p' are mapped into the concave part of the diagram. For this we notice that

$$\frac{\partial J_2(t,u)}{\partial u} = \mathbf{N}_{ijkl}(\mathbf{q}_0) \tag{29}$$

for $t > 1-\frac{1}{2}u$, while $J_2(t,u) \in \text{conv}(V)$ if $t \leq u$. Therefore for $t > 1-\frac{1}{2}u$ the following is true:

$$\frac{\partial C_{i'j'k'l'}(J_2(t_1,u))}{\partial u} \leq 1 \tag{30}$$

$$\frac{\partial C_{ijkl}(J_2(t_1,u))}{\partial u} = 1$$

So as u increases, all constraints increase no faster than C_{ijkl}. So another constraint can only become critical if its lowest common ancestor with C_{ijkl} is a min node, and such a configuration must be in conc(V).

For continuity of J_3, first we notice that J_2 is continuous, and J_3 will be continuous if $u_c(t)$ defined in (28) is continuous. For this we notice that the use of u_v in the min function in (28) guarantees that C_{ijkl} is critical at configuration $J_3(t,u)$ for all t and u. The rest of the proof of continuity is identical to the proof of continuity of $u_c(t)$ in lemma 4.1, using the rate of change condition in (30). □

Lemma 4.4 *If $p : I \to F$ is a path having $p([0,\frac{1}{2}]) = \{p(0)\} \subset \text{conv}(V)$ then p is homotopic to a path $p' : I \to V$ such that $p'(0) = p(0)$.*

Proof. The proof is inductive. We define a sequence of partial homotopies, that is, maps $J^n : I \times [u^{n-1}, u^n] \to F$ such that $J^n(t,u^n) = J^{n+1}(t,u^n)$. Then we show that the number m of partial homotopies required is finite. (We employ a superscript notation which will prove convenient in our inductive argument).

Inductive hypothesis. The input to our construction is a path p^n, a value of $u^n \in I$, and points t_0^n and t_1^n in I such that

(i) $p^n(t) \in \text{conv}(V)$ for $t \leq t_0^n$

(ii) $p^n([t_0^n, t_1^n]) = \{p^n(t_0^n)\} \subset \text{conv}(V)$

(iii) C_{ijkl} is critical on $p^n(I)$.

From C_{ijkl} we use lemma 4.1 to define a homotopy J_1 for the path segment $p^n|_{[t_1^n,1]}$ (reparametrized to I). Similarly, let $C_{i'j'k'l'}$ be a constraint which is critical at $p^n(t_0^n)$ and whose lowest common ancestor with C_{ijkl} is a max node. For these two constraints, we use lemma 4.3 to define an unravelling homotopy J_3 of the path segment $p^n|_{[t_0^n,t_1^n]}$ (reparametrized to I). This gives us a value of u_v as in (27), and we define $u^{n+1} = u^n + u_v$. Now J^{n+1} can be defined on the range $u \in [u^n, u^{n+1}]$ as:

$$J^{n+1} = \begin{cases} p^n(t) & \text{for } t \in [0, t_0^n]; \\ J_3\left(\frac{t-t_0^n}{t_1^n-t_0^n}, u-u^n\right) & \text{for } t \in [t_0^n, t_1^n]; \\ J_1\left(\frac{t-t_1^n}{1-t_1^n}, u-u^n\right) & \text{for } t \in [t_1^n, 1]; \end{cases} \tag{31}$$

then J^{n+1} is a homotopy because J_3 and J_1 agree on their intersection, and $J_3(0,u) = p^n(t_0^n)$. We can define a new path $p^{n+1}(t) = J^{n+1}(t, u^{n+1})$ and points

$$t_0^{n+1} = t_0^n + \frac{u_v}{2}(t_1^n - t_0^n) \quad \text{and} \quad t_1^{n+1} = t_1^n - \frac{u_v}{2}(t_1^n - t_0^n) \tag{32}$$

and it is readily verified that these satisfy the inductive hypothesis.

For the base case, we set $p^0 = p$, $u^0 = 0$ and $t_0^0 = 0$, $t_1^0 = \frac{1}{2}$, which clearly satisfies the inductive hypothesis.

Finiteness. For termination, we must show that after a finite number of steps m, $p^m(t_0^m) \in \text{conc}(V)$. Suppose $p^n(t_0^n) \notin \text{conc}(V)$, and let $C_{i'j'k'l'}$ be the constraint (along with C_{ijkl}) which is critical at $p^n([t_0^{n-1}, t_0^n])$, i.e. these are the constraints used to define the homotopy J_3 for J^n. A third constraint $C_{i''j''k''l''}$ is also critical at $p^n(t_0^n)$, by (27). This constraint has a max node lowest common ancestor with $C_{i'j'k'l'}$, which implies that

$$\frac{\partial}{\partial u} C_{i'j'k'l'}(J_3(\tfrac{1}{2}, u)) < \frac{\partial}{\partial u} C_{i''j''k''l''}(J_3(\tfrac{1}{2}, u)) \tag{33}$$

and from (24) it follows that

$$(N_{ijkl} + N_{i'j'k'l'}) \cdot N_{i'j'k'l'} < (N_{ijkl} + N_{i'j'k'l'}) \cdot N_{i''j''k''l''} \tag{34}$$

which implies

$$N_{ijkl} \cdot N_{i'j'k'l'} < N_{ijkl} \cdot N_{i''j''k''l''} \tag{35}$$

but the condition (35) defines a total ordering on the constraints distinct from C_{ijkl}. That is, as u increases and $p^n(t_0^n)$ is deformed according to some J_2, a new constraint $C_{i''j''k''l''}$ can only become critical if it satisfies condition (35). Once $C_{i''j''k''l''}$ has become critical, $C_{i'j'k'l'}$ can never again become critical. Thus we need define homotopies J^n at most once for each constraint, and so their number m is bounded by the number of constraints.

We then construct J_1 for the path segment $p^m|_{[t_1^m,1]}$ reparameterized to I. The final homotopy is defined for the range $u \in [u^m, 1]$ as:

$$J^{m+1}(t,u) = \begin{cases} p^m(t) & \text{for } t \in [0, t_1^m]; \\ J_1(\frac{t-t_1^m}{1-t_1^m}, u-u^m) & \text{for } t \in [t_1^m, 1]; \end{cases} \tag{36}$$

then the homotopy

$$J_4(t,u) = J^n(t,u) \text{ with } u \in [u^{n-1}, u^n] \text{ for } n = 1, \ldots m+1 \tag{37}$$

is continuous, and defines a homotopy between $p(t) = J_4(t,0)$ and a path $p'(t) = J_4(t,1)$ such that $p'(I) \subset V$. □

Theorem 4.5 *Let $p : I \to F$ be a path with endpoints in V. Then p is path homotopic in F to a path p' with the same endpoints which lies entirely in V.*

Proof. We first apply the homotopy of lemma 4.2 to all path segments $p|_{[t_n,t_{n+1}]}$ with an endpoint in conv(V). This does not displace endpoints in V. Then we construct a global homotopy J by pasting together homotopies J_1 for path segments with an endpoint in conc(V), and homotopies J_4 for the remaining path segments. The resulting homotopy is continuous if these homotopies agree at their endpoints. Firstly, both J_1 and J_4 to not displace endpoints in V. Therefore they agree at endpoints in V. At free endpoints both satisfy the *free endpoint condition*: Assume that after reparametrization, $p(1)$ is a free endpoint. Then

$$J_n(1,u) = p(1) + \min(u, u_c)\mathbf{N}_{ijkl} \qquad \text{for} \qquad n = 1,4 \tag{38}$$

$$\text{with} \qquad u_c \;=\; \inf\{u|\; p(1) + u\mathbf{N}_{ijkl} \in \text{conc}(V)\}$$

thus J is continuous, and we define $p'(t) = J(t,1)$. Since $J_1(I,1) \subset V$ and $J_4(I,1) \subset V$, we have $p'(I) \subset V$. □

Finally, suppose that in a motion planning problem we are given a start configuration $(\mathbf{x},\mathbf{q})$ which is not on V. Then exactly one constraint C_{ijkl} is critical there. We apply the homotopy J_1 to the constant path at $(\mathbf{x},\mathbf{q})$ to attain the diagram; that is, we plan a straight-line path in direction $\mathbf{N}_{ijkl}(\mathbf{q})$ to reach V from the start.

The completeness condition for motion planning has the following simple algebraic formulation. Let $i : V \hookrightarrow F$ be the inclusion map. Then if V is a euclidean Voronoi diagram, then it is a strong deformation retract of F, and hence i induces an isomorphism of fundamental groups. In our case we have the weaker completeness condition that i induces an epimorphism:

Corollary 4.6 *(Algebraic formulation of the completness condition for motion planning). Let $i : V \hookrightarrow F$ be the inclusion map of the simplified Voronoi diagram in free space, with $y_0 \in V$, and let $\pi_1(X,x)$ denote the fundamental group of X with base point x. Then the induced homomorphism $i_* : \pi_1(V,y_0) \to \pi_1(F,y_0)$ is surjective.*

Hence the fundamental group of F is isomorphic to the quotient group $\pi_1(V,y_0)/\ker i_*$. This quotient measures the structural difference between F and V.

5 Complexity bounds

We have given a definition of the simplified Voronoi diagram V in the configuration space of a polyhedron in 3-space. This definition does not constitute an algorithm, so our bounds depend on the algorithm used with the diagram. We assume that the diagram will be used as input to a version of the roadmap algorithm of Canny (1986a). This algorithm computes

one-dimensional skeletons of semi-algebraic sets in time $(d^{O(r^2)}n^r \log n)$ for a semi algebraic set defined by n polynomials of degree d in r variables. In our case the number of variables and the degree of the equations are constants.

A naive bound on the complexity of computing a skeleton of V would be $O(n^{12} \log n)$ if we are given n constraints, because the diagram is a subset of the zero sets of all $O(n^2)$ bisectors of constraints. This bound can be reduced to $O(n^7 \log n)$ by noticing that the diagram has a simple stratification (decomposition into a union of disjoint manifolds). The diagram is a subset of the set of all m-sectors, where an m-sector is the set of points where m constraints have the same value. If the constraints are in general position, each m-sector is a manifold of codimension $m-1$. There are $O(n^m)$ m-sectors, and by the codimension condition, m must be less than or equal to 7. The complexity of computing the skeleton of this stratification is $O(n^7 \log n)$.

While its worst case bounds are poor, the actual performance of the algorithm is expected to be much better, because V approximates the euclidean Voronoi diagram, as shown in figures 4 and 5. The evidence for this is that the complexity of the euclidean Voronoi diagram for a set of n points in r dimensions is $O(n^{\lfloor \frac{(r+1)}{2} \rfloor})$, and the euclidean Voronoi diagram for disjoint line segments in the plane has linear size.

This conjecture is supported by some experimental evidence. We have implemented an algorithm for constructing the simplified Voronoi diagram for the following configuration spaces: $\Re^2$, the case of an arbitrary polygon translating in the plane amidst polygonal obstacles, and $\Re^2 \times S^1$, which allows the moving polygon to rotate as well as translate. In many cases the size of V has been observed to remain roughly linear, as in fig. 4, which our implementation produced.

References

1. Canny J. F., "Collision Detection for Moving Polyhedra," IEEE trans. PAMI, vol 8, no 2, March 1986.

2. Canny J. F., "A Voronoi Method for the Piano-Movers Problem," Proc. IEEE Int. conf. Robotics and Automation, March 1985.

3. Canny, J. F., "Constructing Roadmaps of Semi-Algebraic Sets," Intl. Workshop on Geometric Reasoning, Oxford University, England, June, 1986a.

4. Donald, B. R., "Motion Planning with Six Degrees of Freedom," MIT AI-TR-791, MIT Artificial Intelligence Lab., 1984.

5. Lee, D. T., and Drysdale, R. L., "Generalization of Voronoi diagrams in the plane," SIAM J. Comp. (10) (1981) pp. 73-87.

6. Lozano-Pérez T., "Spatial Planning: A Configuration Space Approach," IEEE Trans. Computers, C-32, No 2 (Feb 1983) pp 108-120.

7. Lozano-Pérez T., and Wesley M., "An algorithm for Planning Collision-Free Paths Among Polyhedral Obstacles", Comm. ACM, vol 22, no 10, Oct 1979.

8. Ó'Dúnlaing, C., Sharir, M., and Yap C., "Generalized Voronoi diagrams for moving a ladder: I Topological Analysis," NYU-Courant Institute, Robotics Lab. Tech. report No. 32 (1984)

9. Ó'Dúnlaing, C., Sharir, M., and Yap C., "Generalized Voronoi diagrams for moving a ladder: II Efficient construction of the diagram," NYU-Courant Institute, Robotics Lab. Tech. report No. 33 (1984)

10. Ó'Dúnlaing C., and Yap C., "A retraction method for planning the motion of a disc," J. Algorithms (6) (1985) pp. 104-111

11. Schwartz J. and Sharir M., "On the 'Piano Movers' Problem, II. General Techniques for Computing Topological Properties of Real Algebraic Manifolds," Comp. Sci. Dept., New York University report 41, 1982.

12. Schwartz J. and Yap C.K., "Advances in Robotics," Lawrence Erlbaum associates, Hillside New Jersey, 1986.

13. Whitney, H., "Elementary Structure of Real Algebraic Varieties," Ann. Math. (66) 3, 1957.

14. Yap, C., "Coordinating the motion of several discs," NYU-Courant Institute, Robotics Lab. No. 16 (1984)

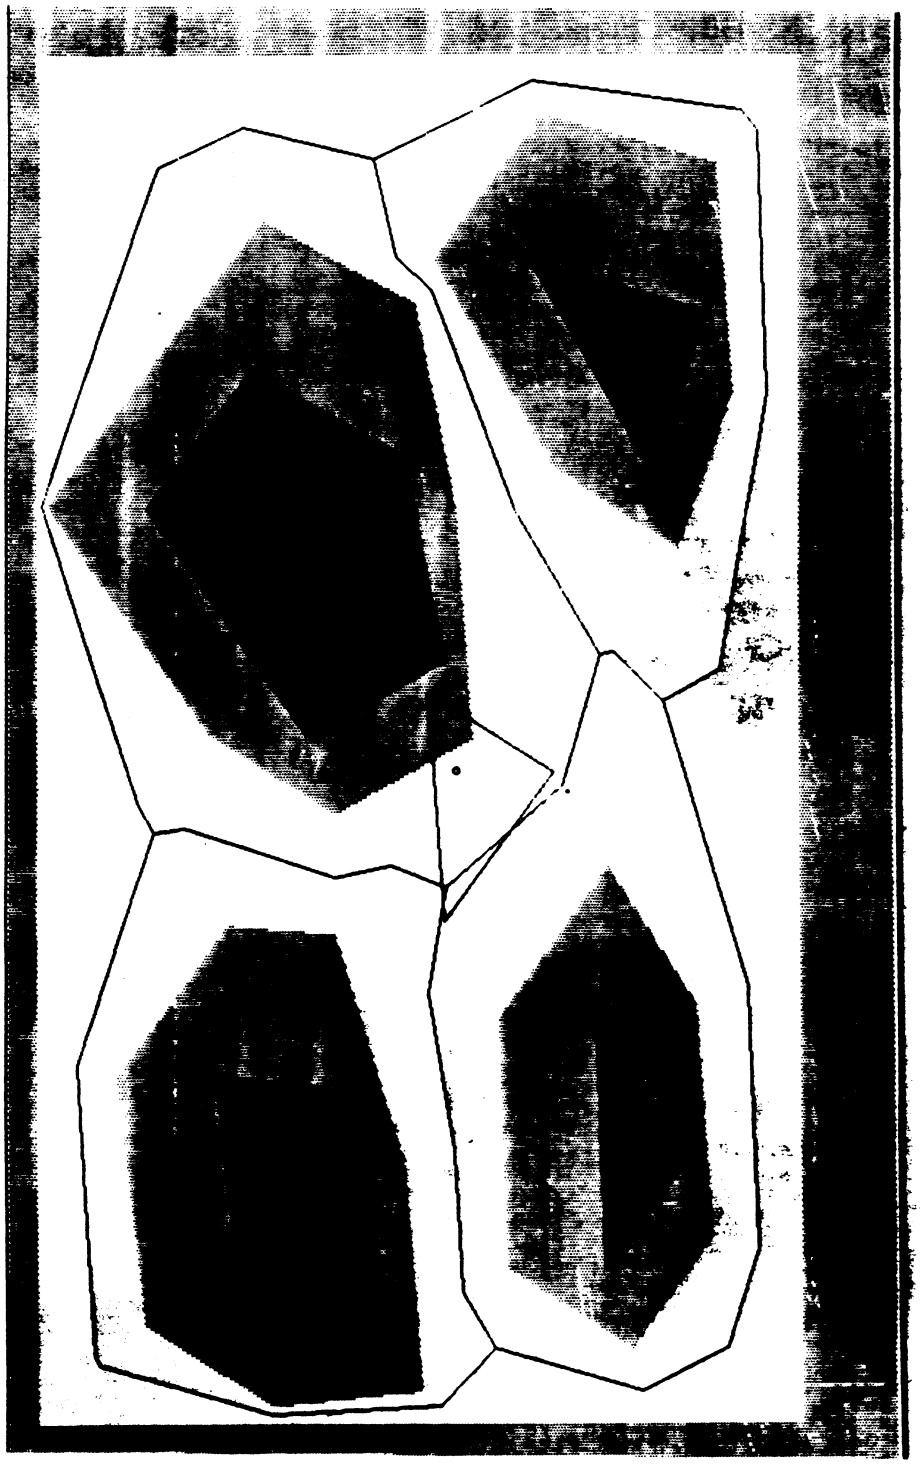

Figure 1. The simplified voronoi diagram in the plane. The black polygons are the real space obstacles. The triangle is the moving "robot." The shaded polygons are the configuration space obstacles. The simplified voronoi diagram is a network of straight line segments

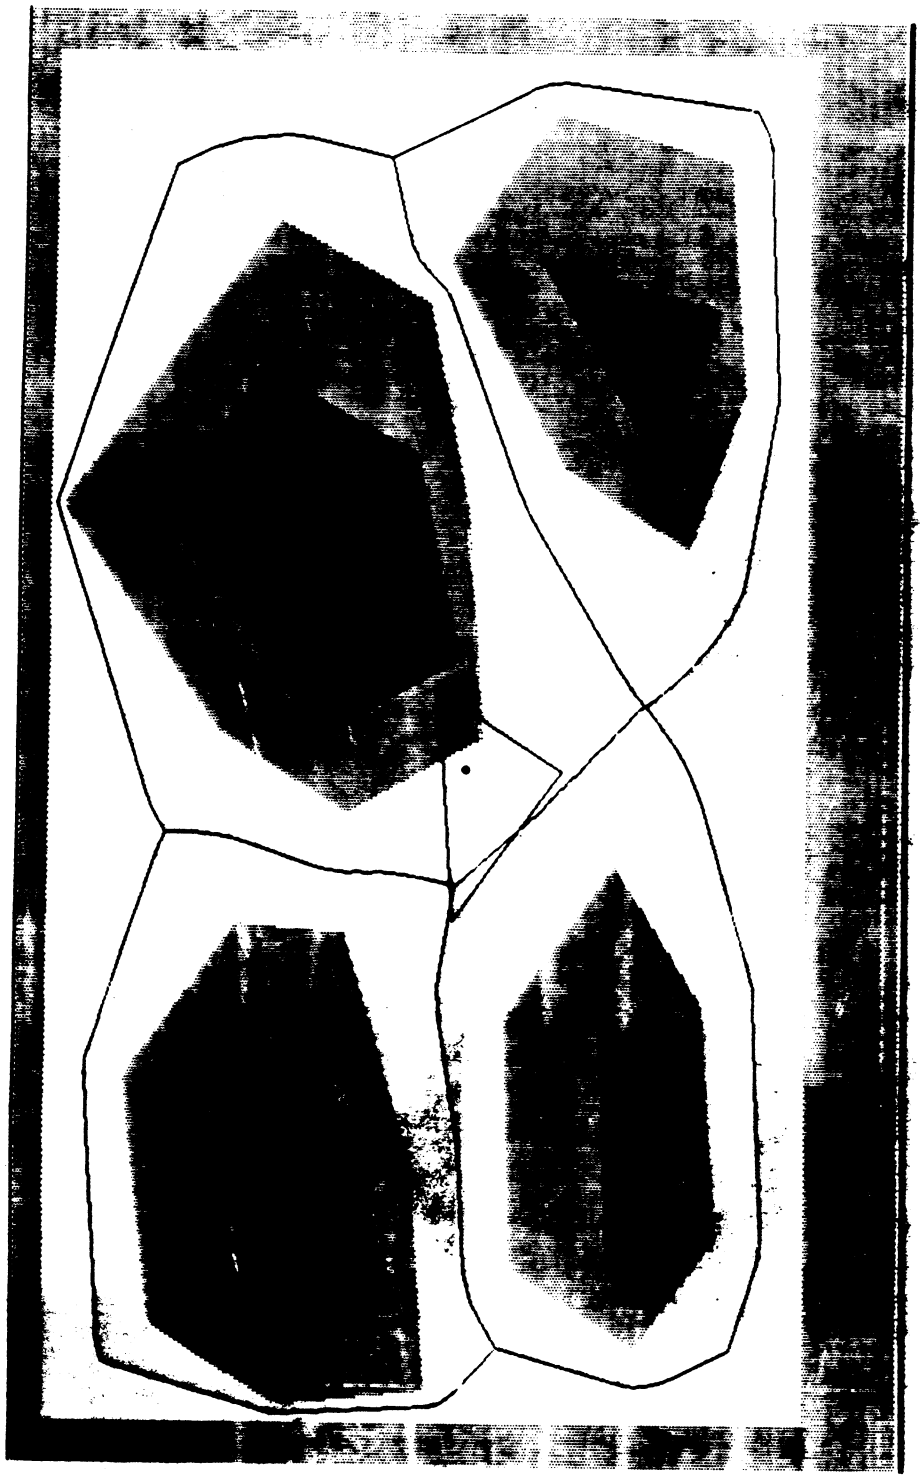

Figure 2. True voronoi diagram for the same obstacles.