

MAPRM: A Probabilistic Roadmap Planner with Sampling on the Medial Axis of the Free Space

Steven A. Wilmarth* Nancy M. Amato† Peter F. Stiller‡

Texas A&M University, College Station, TX 77840

Abstract

Probabilistic roadmap planning methods have been shown to perform well in a number of practical situations, but their performance degrades when paths are required to pass through narrow passages in the free space. We propose a new method of sampling the configuration space in which randomly generated configurations, free or not, are retracted onto the medial axis of the free space. We give algorithms that perform this retraction while avoiding explicit computation of the medial axis, and we show that sampling and retracting in this manner increases the number of nodes found in small volume corridors in a way that is independent of the volume of the corridor and depends only on the characteristics of the obstacles bounding it. Theoretical and experimental results are given to show that this improves performance on problems requiring traversal of narrow passages.

1 Introduction

Motion planning in the presence of obstacles is an important problem in robotics with applications in other areas, such as simulation and computer aided design. While complete motion planning algorithms do exist, they are rarely used in practice since they are computationally infeasible in all but the simplest cases. For this reason, recent attention has focused on probabilistic methods, which sacrifice completeness in favor of computational feasibility and applicability. In particular, several algorithms, known collectively as *probabilistic roadmap planners*, have been shown to perform well in a number of practical situations, see, e.g., [7]. The idea behind these methods is to create a graph of randomly generated collision-free configurations with connections between these nodes made by a simple and fast local planning method. Actual global planning is then carried out on this graph. These methods run quickly and are easy to implement; unfortunately there are simple situations in which they perform poorly, in particular situations in which paths are required to pass through narrow

passages in configuration space.

The *medial axis* or *generalized Voronoi diagram* has a long history of use in motion planning, see [2, 8]. This stems from the fact the medial axis $MA(F)$ of the free space F (the set of all collision-free configurations) has lower dimension than F but is still a complete representation for motion planning purposes. In particular, in two dimensions the medial axis is a one dimensional graph-like structure which can be used as a roadmap. Paths on the medial axis also have other appealing properties such as large clearance from obstacles. However, the medial axis is difficult and expensive to compute explicitly, particularly in higher dimensions.

We propose a new algorithm, MAPRM, which combines these two approaches by generating random networks whose nodes lie on the medial axis of the free space. Our central observation is that *it is possible to efficiently retract any sampled configuration, free or not, onto the medial axis of the free space without having to compute the medial axis explicitly*. Sampling and retracting in this way will be shown to give improved performance on problems requiring traversal of narrow passages.

1.1 Probabilistic roadmap methods

Probabilistic roadmap methods generally operate as follows, see, e.g., [7]. During a preprocessing phase, a set of configurations in the free space is generated by sampling configurations at random and removing those that put the workpiece in collision with an obstacle. These nodes are then connected into a roadmap graph by inserting edges between configurations if they can be connected by a simple and fast local planning method, e.g., a straight line planner. This roadmap can then be queried by connecting given start and goal configurations to nodes in the roadmap (again using the local planner) and then searching for a path in the roadmap connecting these nodes. Various sampling schemes and local planners have been used, see [1, 6, 11]. The algorithms are easy to implement, run quickly, and are applicable to a wide variety of robots.

The main shortcoming of these methods is their poor performance on problems requiring paths that pass through narrow passages in the free space. This is a direct consequence of how the nodes are sampled from F . For example, using the usual uniform sampling over F , any corridor of sufficiently small volume is unlikely to contain any sampled nodes whatsoever. Some effort has been made to

*Department of Mathematics. Supported in part by NSF Group Infrastructure Grant DMS 96-32028.

†Department of Computer Science. Supported in part by NSF CAREER Award CCR-9624315 (with REU Supplement), NSF Grants IIS-9619850 (with REU Supplement), EIA-9805823, and EIA-9810937, and by the Texas Higher Education Coordinating Board under grant ARP-036327-017.

‡Department of Mathematics. Supported in part by the Air Force Office of Scientific Research.

modify sampling to increase the number of nodes sampled in narrow corridors. Intuitively, such narrow corridors may be characterized by their large surface area to volume ratio: the methods in [1] and [5] have exploited this idea.

In [1], nodes are sampled from the *contact space*, the set of configurations for which the workpiece is in contact (but not collision), with an obstacle. This method has solved some very difficult problems, however it is difficult to analyze its performance because the sampling distribution is unknown.

In [5], preliminary configurations are generated by allowing the workpiece to penetrate the obstacles by a small amount. The areas near these nodes are then resampled to find nearby collision-free configurations. Again the idea is that the allowed penetration dilates the free space by a small amount (albeit not uniformly), and the sampling in a narrow corridor is increased roughly in proportion to the surface area. As the authors point out, dilating the free space may alter its topology, opening corridors where none existed. In practice, the amount of dilation must be carefully regulated to mitigate this effect.

1.2 Our results

In this paper we present MAPRM, a new sampling scheme which retracts sampled nodes onto the medial axis of the free space prior to their connection to form a roadmap. The key results are:

- It is possible to efficiently retract almost any configuration, free or not, onto the medial axis of the free space without having to compute the medial axis explicitly.
- Sampling and retracting in this manner increases the number of nodes found in narrow (small volume) corridors in a way that is independent of the volume of the corridor and depends only on the characteristics of the obstacles bounding it.
- This improves performance on problems requiring traversal of such corridors.

A typical approach using the medial axis in motion planning is to compute the medial axis of the free space, which has lower dimension, and to carry out the planning there instead. This is valid because $MA(F)$ is a *strong deformation retract* (SDR) of F , meaning that F can be continuously deformed onto $MA(F)$ while maintaining its topological structure. In fact, as we will show, almost the entire configuration space, free and collision configurations alike, can be retracted onto $MA(F)$. Now, although a complete representation of the medial axis of the free space is difficult and costly to compute, the final retracted image on $MA(F)$ of a given free configuration can be computed efficiently without such a representation. We exploit this fact by sampling nodes from the full configuration space and retracting them onto $MA(F)$. These nodes, now all

in the free space, can be connected in the usual way to form a roadmap. We will show that this has the effect of increasing sampling in narrow corridors in a way that is independent of the volume of the corridor.

We give a theoretical treatment and experimental results for two dimensional configuration space, and present the algorithm and experimental results for the case of a free-flying rigid body in three dimensions. A theoretical treatment of the rigid body case will be presented in the full paper.

2 Retracting onto the medial axis in the plane

In this section we give the theoretical development for the plane, which we interpret as a two dimensional configuration space. We first define the medial axis of a polygon in the plane and show that it is a strong deformation retract; this essentially means that the polygon can be continuously deformed onto its medial axis while maintaining its topological structure. In the context of motion planning, if we consider the free space F to be a polygon (with holes), this implies that it is valid to restrict attention to the medial axis $MA(F)$, in the sense that there is a path between two configurations in F if and only if there is a (homotopic) path between their images on $MA(F)$. We will then show how this map may be extended to retract almost all of the full configuration space onto $MA(F)$. The MAPRM algorithm, uniform sampling followed by application of this extended retraction map, will be shown to yield the desired sampling increase in the narrow passages.

2.1 Retracting a polygon onto its medial axis

We momentarily leave the motion planning setting and just consider the medial axis for a region in the plane. In the next section we will regard the plane as a two dimensional configuration space in which the free space and obstacles are polygonal. It is crucial to note that in higher dimensions, we are interested in the medial axis of the free space, which lies in the configuration space and not in the workspace.

Although this development could proceed assuming only piecewise real analytic boundary (see [4]), for simplicity we consider only sets P that are the disjoint union of a finite number of closed polygons (including the interior, possibly with holes). For $x \in P$, we define $B_P(x)$ to be the largest closed disc centered at x that is a subset of P , i.e., $B_P(x) = \overline{B}(x, \rho_P(x))$, where $\overline{B}(x, r)$ denotes the closed disc of radius $r \geq 0$ centered at x , and $\rho_P(x) = \text{dist}(x, \mathbb{R}^2 \setminus P)$ is the distance to the boundary for points inside P , and 0 for points outside P .¹ The *medial axis* $MA(P)$ of P is defined to be the set of all points x of P whose associated $B_P(x)$ are maximal; i.e.,

$$MA(P) = \{x \in P \mid \nexists y \in P \text{ with } B_P(x) \subsetneq B_P(y)\}.$$

¹We define $\text{dist}(x, S) = \inf_{y \in S} d(x, y)$ and $\text{dist}(R, S) = \inf_{x \in R} \text{dist}(x, S)$ where d denotes the Euclidean distance.

Figure 1 shows an example of a polygon and its medial axis. A point $x \in P$ is called a *simple point* if x has a




Figure 1: Medial axis of a polygon

unique nearest point $x_0 \in \partial P$ (so that $d(x, x_0) = \rho_P(x)$). Otherwise, x is called a *multiple point*.

There is a well-defined notion of convexity of P at a vertex. Vertices at which P is convex are in fact always on the medial axis. Other vertices are less convenient to deal with: we define P' to be P minus its non-convex vertex points. We collect a few facts about the medial axis.

Proposition 2.1. *Let P be as above. Then:*

1. *Let $x \in \partial P$. Then x is in $\text{MA}(P)$ if and only if x is a convex vertex of P .*
2. *Any multiple point of P is contained in $\text{MA}(P)$. If $x \in \text{MA}(P)$, then x is in the interior of P if and only if x is a multiple point of P .²*
3. *For each $x \in P'$, $B_P(x)$ is contained in a unique maximal disc $B_P(y)$, where $y \in \text{MA}(P)$. Furthermore if $x \in P^\circ \setminus \text{MA}(P)$, then y is on the ray $\overrightarrow{x_0 x}$ where $x_0 \in \partial P$ is the unique nearest boundary point to x . If x is a non-vertex point of the boundary, then y is on the line through x normal to the boundary at x .*
4. *The map $P \setminus \text{MA}(P) \rightarrow \partial P$ taking each point to its nearest boundary point is continuous.*
5. *$\text{MA}(P)$ is a closed set.*

Proof. See [13]. □

A subset Y of a set $X \subseteq \mathbb{R}^2$ is called a *strong deformation retract* (SDR) of X if X can be continuously deformed onto Y without moving any of the points of Y ; i.e., there must exist functions $h_t : X \rightarrow X$ for $t \in [0, 1]$ such that: $h_t(y) = y$ for all $y \in Y$ and all $t \in [0, 1]$; $h_0(x) = x$ and $h_1(x) \in Y$ for all $x \in X$; and $h_t(x)$ is a continuous function of t and x . The importance of this concept is that Y retains the topological structure of X ; in particular, there is a path between two points of X if and only if there is a (homotopic) path in Y between their retracted images on Y . If we imagine X to be the free space for

²For more general domains, there may be simple points in $\text{MA}(P) \cap P^\circ$. See [4].

some planning problem, then this would allow us to do the planning on Y instead; in our case we would like to know that $Y = \text{MA}(X)$ is an SDR of X .

We now show that for our polygon P , $\text{MA}(P)$ is in fact an SDR of P' . First, define a map r_P taking P' to $\text{MA}(P)$ by mapping each point x not in $\text{MA}(P)$ to the unique point $y \in \text{MA}(P)$ such that $B_P(x) \subseteq B_P(y)$, as given by Proposition 2.1 Part 3. (See Figure 2.)




Figure 2: Image of a point x under the canonical retraction map.

We will show below that r_P is continuous, so we simply define $h_t(x) = (1-t)x + t r_P(x)$; observe that because the segment connecting x and $r_P(x)$ always lies in P , h_t does in fact map into P . The map r_P is known as the *canonical retraction map*. In Figure 1, the dashed lines show the paths along which points move during the retraction. Note the obvious continuity problems if we try to include vertices at which P is not convex.³

As an immediate consequence we can actually compute the retracted image $r_P(x)$ on $\text{MA}(P)$ of any point x in P' without having to first compute $\text{MA}(P)$. Using a bisection method, for example, we can simply search along the ray $\overrightarrow{x_0 x}$, to find arbitrarily close points x_a and x_b on $\overrightarrow{x_0 x}$, where x_0 is a nearest boundary point for x_a but not for x_b . Then $r_P(x) \in \overline{x_a x_b}$. The only operation required here is the ability to compute the nearest boundary point, which is supported by many collision detection algorithms.

Proposition 2.2. *$\text{MA}(P)$ is an SDR of P' under the map h_t given above.*

Proof. See [13]. □

2.2 Extending the retraction map

We now return to the setting of motion planning for a planar configuration space or C -space. The configuration space can be broken into two pieces which are essentially

³It is true that $\text{MA}(P)$ is an SDR of all of P – showing this simply involves first “shrinking” P slightly into its interior to make it a subset of P' , and then applying r_P .

complements of each other: the free space F of collision free configurations, and the configuration space obstacles or C -obstacles. If we assume that the free space F is a polygon (with holes), we can use the results of the previous section to get a retraction map taking F (actually F') onto its medial axis $MA(F)$; while this may have some benefit such as increasing clearance from obstacles, there is no obvious sampling increase in narrow passages. However, we will show in this section how collision configurations can also be retracted onto $MA(F)$; this extended retraction map is what will give the desired sampling increase in narrow passages. We use the following setting:

- Let C , the *configuration space*, be a closed rectangular region in \mathbb{R}^2 .
- Let $B \subseteq C$, the *C-obstacle*, be a disjoint union of a finite number of polygons.
- $F = \overline{C \setminus B}$ is called the *free space*. F consists of a finite number of connected components, each of which is a polygon.
- The set ∂F is called the *contact space*. For simplicity, we assume that ∂B is disjoint from ∂C , so that $\partial B \subseteq \partial F$.

As in the previous section, we have the canonical retraction map r_F taking F' onto its medial axis $MA(F)$. We now show how this map may be extended to take $C \setminus MA(B)$ onto $MA(F)$. The idea is simple: retraction of points of F' to $MA(F)$ moves points along lines that meet the non-vertex points of the boundary of F normally. These lines can be continuously extended into B until they meet the medial axis of B , so that any point of $B \setminus MA(B)$ can be moved toward F along these lines. This amounts to moving each point of B toward its nearest boundary point on ∂B and through ∂B into F until $MA(F)$ is reached. See Figure 3.




Figure 3: The canonical retraction map (a) and extended retraction map (b). The shaded area is B .

Proposition 2.3. *With C , B , and F as above, the canonical retraction map $r_F : F' \rightarrow MA(F)$ can be extended continuously to map $C \setminus MA(B) \rightarrow MA(F)$.*

Proof. By Proposition 2.1 part 4, we can map $B \setminus MA(B)$ continuously onto the boundary $\partial B \subseteq \partial F$. Any non-convex vertex of F is a convex vertex of B ; because a convex vertex of B is already in $MA(B)$ and is not a nearest boundary point for any point of $B \setminus MA(B)$, we have in fact mapped $B \setminus MA(B)$ into F' . We may then apply the canonical retraction map r_F . \square

We call this map the *extended retraction map*. Using the same definition of h_t as before, it is easy to show that in fact $MA(F)$ is an SDR of $C \setminus MA(B)$.

Again we can compute this extended retraction map for a given $x \in C \setminus MA(B)$ without first computing $MA(F)$. Let x_0 be the nearest point on ∂B to x . If x is in F , move x away from x_0 along the ray $\overline{x_0 x}$ as before. If x is in B , move away from x_0 along the ray $\overline{x x_0}$ starting at x_0 . (If x is actually on ∂B , we move in the direction given by the outward normal to ∂B at x .) The MAPRM algorithm, listed in detail below (Algorithm 3.1), is just uniform sampling followed by application of this map. Note that in the case of polygons in the plane, the medial axis $MA(B)$ has measure zero, see [4], so in practical terms we are able to retract any configuration sampled from C .

2.3 Sampling is increased in narrow corridors

We next examine the sampling distribution obtained by MAPRM in sampling uniformly from C and applying the extended retraction map. We will show that this sample-and-retract scheme improves sampling in small volume corridors.

The medial axis provides a convenient definition for what is meant by a corridor. Let $r_F : F' \rightarrow MA(F)$ be the canonical retraction map. A *corridor* in F is a connected subset S of F' such that $r_F(S) \subseteq S$ and $r_F^{-1}(S \cap MA(F)) \subseteq S$. These conditions essentially ensure that S is bounded on “both” sides by obstacles.

We can easily compute a lower bound on the volume of points that must map into such a corridor under the extended retraction map. Clearly any point of such a corridor S remains in S under the extended retraction map. Furthermore, any point in B whose nearest boundary point (on ∂B) is also in S will be mapped into S ; if $x \in \partial S \cap B'$, then x is the nearest boundary point for all points on the (open) segment $\overline{x r_B(x)}$. These segments are normal to the piecewise linear boundary of B' , so we have:

Proposition 2.4. *Let $S \subset F'$ be a corridor, $r_B : B' \rightarrow MA(B)$ the canonical retraction map for B , and $q : C \setminus MA(B) \rightarrow MA(F)$ the extended retraction map. The volume of points $x \in C$ that map into S under q is no smaller than:*

$$\text{Vol}(S) + \int_{\partial S \cap B'} d(x, r_B(x)) dl(x). \quad (2.1)$$

(dl is the unit of arc length along the boundary.)

Note that although we have increased the volume of points that produce nodes in the corridor (as compared to uniform sampling), we are also sampling from a larger set: all of C rather than just F . So the sampling is not improved in *every* corridor, but only in those with sufficiently small volume.

We can view the function $d(x, r_B(x))$ defined on the surface of the obstacle as a measure of how “thick” the obstacle is near x . If $d(x, r_B(x))$ is small, the medial axis is very near the surface and there are other features (edges) nearby. Note that the value of the integral in the above expression depends only on the values of this function on the obstacle surfaces that form the corridor and not on the volume of the corridor itself.

3 MAPRM: Sampling from the medial axis

In this section, we summarize the MAPRM algorithm in two dimensions and give some experimental results.

3.1 The MAPRM algorithm in 2D

The MAPRM algorithm for construction of a roadmap consists of uniform sampling in the the full configuration space, followed by application of the extended retraction map. The nodes are then connected to form a roadmap using a local planner as usual.

MAPRM in the plane is given in Algorithm 3.1.

Algorithm 3.1 MAPRM in 2D

Preprocessing:

Input. N , the number of nodes to generate.

Output. N nodes in F connected into a roadmap graph.

- 1: **repeat**
 - 2: Generate a uniformly random point p in C .
 - 3: Find the nearest point q on ∂F to p .
 - 4: **if** p is free **then**
 - 5: Take the retraction direction \vec{v} to be $\vec{q\bar{p}}$, and let the start point s be p .
 - 6: **else**
 - 7: Take the retraction direction \vec{v} to be $\vec{p\bar{q}}$, and let the start point s be q .
 - 8: **end if**
 - 9: Using bisection, move s in the direction \vec{v} until q is not the unique nearest point of ∂F to s . This moves s onto the medial axis of F
 - 10: **until** N nodes have been generated
 - 11: For each pair of nodes: if the pair can be connected with a straight line, insert an edge into the graph connecting them.
-

3.2 Examples

We give two examples: one in which MAPRM shows significant advantage over uniform sampling, and one in which it does not. We omit the connections between nodes and show only the results of the sampling.

Figure 4 shows an example of a free space containing a narrow corridor. Part (a) shows the result of sampling 100 nodes from the free space: this required generating

168 random configurations. Part (b) shows the result of sampling 100 nodes using the MAPRM Algorithm 3.1: 100 points were generated in the square and the extended retraction map was applied to them. MAPRM produces many nodes in the corridor because the obstacles forming the corridor are “thick” which gives a reasonably large value for the integral in Equation (2.1). Because of this, pushing the two obstacles closer together would not greatly affect the number of nodes MAPRM generates in the corridor.




Figure 4: Uniform sampling (a) vs. MAPRM sampling (b)




Figure 5: Uniform sampling (a) vs. MAPRM sampling (b)

Figure 5 shows an example in which MAPRM is less helpful. As before, nodes in (a) are 100 free configurations (203 samples required), and nodes in (b) are 100 samples from the square retracted to $MA(F)$. The medial axis of the obstacles forming the corridor has a “maze-like” quality. The medial axis is very near the surface causing the integral in Equation (2.1) to be smaller in this case.

The running time of both algorithms on these planar examples is insignificant. However, as we will see below in the 3d rigid body case, as the volume of the free space gets smaller, MAPRM actually takes significantly *less* time to generate a free node than uniform sampling, decreasing the overall running time.

4 MAPRM for a rigid body

In this section we explain the extension of MAPRM to the case of a rigid polyhedron moving among polyhedral

obstacles in \mathbb{R}^3 . The goal in this setting is to perform the retraction in configuration space, which in this case is six dimensional. Specifically, we want to retract all configurations, free or otherwise, onto the medial axis of the free space. While avoiding explicit computation of the medial axis as before, we also seek to avoid the costly computation of the C -obstacle boundaries in configuration space. Without explicit knowledge of these boundaries, we are able to retract all free configurations but only a subset of the collision configurations.

The collision configurations that are discarded are those that retract to contact configurations with more than one contact point, i.e., contact configurations that are on the medial axis of the free space. Such contact configurations roughly correspond to the convex vertices of F in the planar case. However, this reduced subset is still sufficient to increase sampling in narrow corridors.

We give only the algorithm in this paper; the theoretical treatment of the rigid body case will be presented in the full paper. The theory is complicated by the issue of choice of metric on the configuration space: the definition of the medial axis depends on the choice of metric, the retraction maps depend on a notion of a line or shortest path between configurations, and even the idea of *uniform* sampling from the configuration space depends on the choice of metric. Our approach is to use a Riemannian metric to impose the necessary geometric structure, but there are still many possible choices, see [12]. Rather than restricting our attention to a particular Riemannian metric, we enumerate our assumptions about the metric being used and give a retraction algorithm which will always work under these conditions. We present the resulting algorithm here, described in terms of geometry and motions in the workspace.

Finally, we give experimental results showing how a preliminary implementation performs on a specific problem requiring traversal of a narrow corridor.

4.1 Motion planning for a rigid body

The configuration space for a free-flying rigid body U in \mathbb{R}^3 describes all possible positions and orientations of U ignoring any obstacles that may be present. A particular configuration of a rigid body may be described by specifying the position and orientation of a moving coordinate system attached to U , the *body frame*, with respect to a particular fixed system, the *world frame*. Such coordinate systems are related by a rotation matrix⁴ in $SO(3)$ giving the orientation of the body frame with respect to the world frame, together with a vector in \mathbb{R}^3 specifying the location of the origin of the body frame with respect to the world frame. We denote the set of all such pairs by $SE(3) = SO(3) \times \mathbb{R}^3$. A particular pair $(R, p) \in SE(3)$ operates on the body frame coordinates q_a of a point to produce the world frame coordinates $Rq_a + p$ of that same

⁴ $SO(3) = \{R \in \mathbb{R}^{3 \times 3} \mid RR^T = I \text{ and } \det(R) = 1\}$

point. If $c = (R, p)$, we write $c \cdot U$ or $(R, p) \cdot U$ to mean the coordinates of all points of U with respect to the world frame when U is in configuration c . See [10] for more detail on $SE(3)$.

Now, given an obstacle V in the workspace \mathbb{R}^3 , certain configurations of U are prohibited because they cause U to overlap V , i.e., $\{c \in SE(3) \mid (c \cdot U) \cap V \neq \emptyset\}$. We call this set of configurations the C -obstacle of V . This divides the configuration space into two pieces: F , the *free space* of collision free configurations, and B the union of the C -obstacles associated with any obstacles present in the workspace.

- We take the configuration space C to be a “rectangular” subset of $SE(3)$; i.e., $C = SO(3) \times W$ where W is a closed rectangular region in \mathbb{R}^3 , which will be convenient for sampling.
- We assume the *workpiece* U and the *obstacles* $V_j, j = 0, \dots, n$ are closed polyhedra in \mathbb{R}^3 (including the interior, possibly with holes), with the V_j (pairwise) disjoint. We let V be the union of the V_j .
- We let the free space F be the closure of C minus the C -obstacles of the V_j .
- Configurations in ∂F are called *contact configurations*; they put U in contact with either V or the boundary of the configuration space.

The problem now is to plan a path of configurations in F between given start and goal configurations in F . We would in particular like to apply the retraction method of the previous section replacing the plane and its polygonal obstacles by $SE(3)$ and its C -obstacles. In addition, we would like to avoid not only explicit computation of the medial axis, but explicit computation of the C -obstacles as well.

4.2 MAPRM for a rigid body

The relevant points of the theory are:

- Let c be a free configuration. If $c \cdot U$ has two distinct nearest obstacle points, i.e., there are two distinct points $y_1, y_2 \in V$ such that $\text{dist}(V, c \cdot U) = \text{dist}(y_1, c \cdot U) = \text{dist}(y_2, c \cdot U)$, then c is on the medial axis of F .
- Retraction of a free configuration onto the medial axis of F amounts to translating U away from the unique nearest point of V until there are at least two distinct nearest obstacle points: i.e., if c is a free configuration not on the medial axis, and $x \in c \cdot U$ and $y \in V$ are nearest points of $c \cdot U$ and V (so that $d(x, y) = \text{dist}(c \cdot U, V)$) then the retraction map translates U away from y along the line \overleftrightarrow{xy} until there is no unique nearest point of V .

- For a collision configuration c , the (extended) retraction map makes the shortest translation that will free U from collision (so U will then be in contact with V) and continues in that direction as above until there is no unique nearest point to U on V .
- For collision configurations, we only retract configurations whose nearest contact configuration puts U in contact with V in a single point. Contact configurations placing U in contact with V in more than one point are already on the medial axis and are analogous to convex vertices of F in the planar case.

The resulting algorithms are given in Algorithms 4.1 and 4.2. Our simple algorithm for computing the nearest contact configuration for a collision configuration is essentially an exhaustive search for the shortest translation.⁵

Algorithm 4.1 MAPRM for rigid bodies in 3D

Preprocessing:

Input. N , the number of nodes to generate.

Output. N nodes in F connected into a roadmap.

- 1: **repeat**
 - 2: Sample a configuration (R, p) from C .
 - 3: Use Algorithm 4.2 to get the nearest contact configuration (R, q) to (R, p) .
 - 4: **if** (R, p) is free **then**
 - 5: Take the retraction direction \vec{v} to be \overrightarrow{qp} , and let the start point s be p .
 - 6: **else**
 - 7: Take the retraction direction \vec{v} to be \overrightarrow{pq} , and let the start point s be q .
 - 8: **end if**
 - 9: Starting in configuration (R, s) , translate U in the direction \vec{v} until there are two nearest points on V to U . This configuration is on the medial axis of the free space.
 - 10: **until** N vertices have been output
 - 11: For each pair of vertices: if the pair can be connected with the local planner, insert an edge into the graph connecting them.
-

4.3 Implementation and experimental results

We implemented MAPRM for rigid bodies using V-Clip [9] to provide collision detection and closest pair calculations. We used a single local planner: translation with simultaneous rotation about the principal axis of rotation. Normally in practice, connections are not attempted between all pairs of nodes, only between “nearby” nodes according to some metric. In our example, so few nodes are generated that it was feasible to attempt connections between all pairs necessary to determine the components of the roadmap.

Our example is shown in Figure 6. The workpiece is a cube of side length 2; the obstacle is a solid cube of side

⁵Some work has been done on this shortest translation problem, but only for convex polyhedra, see [3].

Algorithm 4.2 Finding the nearest contact configuration

Input. A configuration (R, p) .

Output. A shortest translation (R, q) from configuration (R, p) that puts U in contact (but not collision) with an obstacle, or **failure**.

- 1: **if** (R, p) is a free configuration **then**
 - 2: Return $(R, p + (x - y))$, where $x \in (R, p) \cdot U$ and $y \in V$ are a pair of closest points between U and V .
 - 3: **else**
 - 4: For each feature (vertex, edge, face) of U , and each feature of the V_j , find the configuration (R, q) with smallest $|p - q|$ that puts these features in contact in a single point.
 - 5: Such a configuration may put U strictly in collision (not just contact) with an obstacle; discard any such configurations.
 - 6: If no configurations remain, output **failure**.
 - 7: Otherwise, output a remaining configuration (R, q) with smallest $|p - q|$.
 - 8: **end if**
-

length 20 with the indicated corridor cut through it. The corridor has 2.5×2.5 cross section.

We compared experiments using uniform sampling (with collisions discarded) against sampling with MAPRM. The same local planner, collision detection, connection scheme, etc., were used for both methods. Configurations were sampled with arbitrary rotation, and translations placing the center of the workpiece anywhere inside the $20 \times 20 \times 20$ cube. Execution of each method was terminated when some component of the roadmap reached from one mouth of the corridor to the other.

The mean results for 15 runs are given in Table 1.⁶ Observe that on average the MAPRM algorithm solved the problem in less than one-tenth of the time required by uniform random sampling. MAPRM generated nodes at about 13 times the rate of uniform sampling. Note the huge number of random samples required using uniform sampling. However, observe that MAPRM generally required more nodes in the roadmap to be able to solve the problem. We attribute this to the non-uniform distribution of nodes generated by MAPRM along the medial axis: in general there will be somewhat fewer nodes sampled near corners than in the straight sections. However, the much greater sampling rate of MAPRM far outstrips this demand for additional nodes. This effect warrants further investigation.

5 Conclusions

We have described a new sampling method for probabilistic path planning which retracts sampled configurations onto the medial axis of the free space. The method was shown, theoretically and experimentally, to give improved performance on problems requiring traversal of small volume passages in the free space.

Although the method performs well, it employs more

⁶Experiments were run on a MIPS R10000 processor at 200 Mhz.




Figure 6: Rigid body example. The obstacle block is solid with the indicated corridor cut through it.

Table 1: Experimental results (times in seconds)

	MAPRM	Uniform
Sampling time	690	7875
Connection time	82	79
Total preprocessing time	772	7954
# roadmap nodes required	404	351
# random config. sampled	39,568	114,058,889
Mean time to generate a node	1.71	22.65

complicated geometric calculations than the standard uniform sampling and is consequently slightly more difficult to implement. For larger problems, we expect time required for the calculation of the nearest contact configuration to become more significant: additional sophistication will probably be required in that calculation.

In the future, we would like to apply this technique to articulated robots, an area in which probabilistic methods have been very successful. Finally, the existing theoretical results regarding probability of success of the overall probabilistic roadmap method assume uniform sampling of the free space; we would like to prove some similar general results about the overall performance of MAPRM.

References

- [1] N. Amato, O. B. Bayazit, L. K. Dale, C. Jones, and D. Vallejo. OBPRM: An obstacle-based PRM for 3D workspaces. In P. K. Agarwal, L. E. Kavraki, and M. Mason, editors, *Proc. Workshop Algorithmic Found. Robot. A. K. Peters, Wellesley, MA, 1998.*
- [2] F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23:345–405, 1991.
- [3] S. Cameron. Enhancing GJK: Computing minimum and penetration distances between convex polyhedra. In *Intl. Conf. Robotics and Auto.*, 1997.
- [4] H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical theory of medial axis transform. *Pacific J. Math.*, 181(1):57–88, 1997.
- [5] D. Hsu, L. E. Kavraki, J.-C. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In *Proc. 1998 Workshop Algorithmic Found. Robot.*, Wellesley, MA, 1998. A. K. Peters.
- [6] L. Kavraki. *Random Networks in Configuration Space for Fast Path Planning*. PhD thesis, Stanford Univ., Stanford, CA, 1995.
- [7] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high dimensional configuration spaces. *IEEE Trans. Robot. Autom.*, 12:566–580, 1996.
- [8] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.
- [9] B. Mirtich. V-Clip: Fast and robust polyhedral collision detection. Technical Report TR97-05, MERL, 201 Broadway, Cambridge, MA 02139, USA, July 1997.
- [10] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation*. CRC Press, Boca Raton, FL, 1994.
- [11] M. Overmars and P. Svestka. A probabilistic learning approach to motion planning. In *Proc. Workshop on Algorithmic Foundations of Robotics*, pages 19–37, 1994.
- [12] F. C. Park. Distance metrics on the rigid-body motions with applications to mechanism design. *ASME Trans., Journal of Mechanical Design*, 117(1):48–54, 1995.
- [13] S. A. Wilmarth, N. M. Amato, and P. F. Stiller. MAPRM: A probabilistic roadmap planner with sampling on the medial axis of the free space. Technical Report TR98-022, Department of Computer Science, Texas A&M University, College Station, TX, Nov. 1998.