Motion Planning for a Rigid Body Using Random Networks on the Medial Axis of the Free Space

Steven A. Wilmarth*
Dept. of Mathematics
Texas A&M University
wilmarth@math.tamu.edu

Nancy M. Amato†
Dept. of Computer Science
Texas A&M University
amato@cs.tamu.edu

Peter F. Stiller‡
Dept. of Mathematics
Texas A&M University
stiller@math.tamu.edu

Abstract

Several motion planning methods using networks of randomly generated nodes in the free space have been shown to perform well in a number of cases, however their performance degrades when paths are required to pass through narrow passages in the free space. In [16] we proposed MAPRM, a method of sampling the configuration space in which randomly generated configurations, free or not, are retracted onto the medial axis of the free space without having to first compute the medial axis; this was shown to increase sampling in narrow passages. In this paper we give details of the MAPRM algorithm for the case of a free-flying rigid body moving in three dimensions, and show that the retraction may be carried out without explicitly computing the C-obstacles or the medial axis. We give theoretical arguments to show that this improves sampling in narrow corridors, and present preliminary experimental results comparing the performance to uniform random sampling from the free space.

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 give only weaker forms of completeness in exchange for 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., [9]. 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, 10]. This stems from the fact that the medial axis $\mathrm{MA}(F)$ of the free space F (the set of all collision-free configurations) generally has lower dimension than F but is still a complete representation for motion planning purposes. For example, 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.

In [16] we proposed a new algorithm, MAPRM, which combined 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 retract a configuration, free or not, onto the medial axis of the free space without having to compute the medial axis explicitly.* We gave a detailed account for a planar configuration space, showed how sampling and retracting in this way gave improved sampling in narrow passages, and sketched the algorithm for the case of a rigid body moving in 3D. In this paper we present details of the rigid body case and give some analysis and results for this case.

1.1 Probabilistic roadmap methods

Probabilistic roadmap methods generally operate as follows, see, e.g., [9]. 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, 7, 8, 14]. The algorithms are easy to implement, run quickly, and are applicable to a

[1]Supported in part by NSF Group Infrastructure Grant DMS 96-32028.

[2]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.

[3]Supported in part by the Air Force Office of Scientific Research.

Proceedings of the 15th Annual ACM Symposium on Computational Geometry, 1999, pp. 173–180.

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 uniform sampling over F, any corridor of sufficiently small volume is unlikely to contain any sampled nodes whatsoever. Some work has been done on modified sampling strategies aimed at increasing 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 [6] 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 its performance is difficult to analyze because the sampling distribution on the contact space is unknown.

In [6], 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 previous work [16] we introduced MAPRM, a sampling scheme which retracts sampled nodes onto the medial axis of the free space prior to their connection to form a roadmap. We gave a detailed account for a planar configuration space, and sketched the case of a rigid body moving in 3D. In this paper we present details rigid body case and extend the analysis to this case. The key points 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. A consequence of the sampling method is that the sampling is increased more near "thicker" obstacles.
- In the case of a free flying rigid body in 3D, the medial axis in the configuration space has a simple interpretation in the workspace for a large class of configuration space metrics.

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 $\mathrm{MA}(F)$ is a *strong deformation retract* (SDR) of F, meaning that F can be continuously deformed onto $\mathrm{MA}(F)$ 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 $\mathrm{MA}(F)$. Although a complete representation of the medial axis of the free space is difficult and costly to compute, the final retracted image on $\mathrm{MA}(F)$ 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 $\mathrm{MA}(F)$. 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.

In the case of a point moving in the plane, the C-obstacle boundaries are known explicitly and this makes it particularly easy to carry out the retraction. In this paper we show how the retraction map may be carried out in the six dimensional configuration space $\mathrm{SE}(3)$ of a rigid body in 3D without explicitly computing the C-obstacle boundaries.

We give a brief motivation using the planar case, and then describe the details and give analysis and some experimental results for the rigid body case.

2 Preliminaries

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 matrix[1] in $\mathrm{SO}(3)$ 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 $\mathrm{SE}(3) = \mathrm{SO}(3) \times \mathbb{R}^3$. A particular pair $(R, p) \in \mathrm{SE}(3)$ operates on the body frame coordinates q of a point to produce the world frame coordinates $Rq + p$ of that same point.[2] 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 [13] for more detail on $\mathrm{SE}(3)$.

Given an obstacle $V \subset \mathbb{R}^3$ in the workspace, certain configurations of U are prohibited because they cause U to collide with V. 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 *blocked space*, the union of the C-obstacles associated with all obstacles present in the workspace. We use the following setting:

- $W = \mathbb{R}^3$ is the *workspace*.
- The workpiece U is assumed to be a closed and bounded subset of W.
- By taking the union of all obstacles in the workspace, we assume there is a single obstacle V that is a closed subset of W.
- C denotes the configuration space, $\mathrm{SE}(3)$.
- We define $C(U, V)$ to be the set of all configurations of U in C that cause U to meet V, i.e.,

$$C(U, V) = \{g \in C \mid (g \cdot U) \cap V \neq \emptyset\}.$$

$C(U, V)$ is known as the C-obstacle of V.

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

[2] As a group, $\mathrm{SE}(3)$ is the *semi-direct product* $\mathrm{SO}(3) \ltimes \mathbb{R}^3$ where the multiplication is given by composition, i.e., $(R_1, p_1) \cdot (R_2, p_2) = (R_1R_2, R_1p_2 + p_1)$.

- The *blocked space* B is the subset $C(U, V)$ of the configuration space.
- The *free space* of F is the closure of $C \setminus B$.
- Configurations in ∂F are called *contact configurations*: they put U in contact but not overlap with V.

The motion planning problem in this setting is to plan a path of configurations in F between given start and goal configurations in F.

3 MAPRM in the plane

For motivation, we first sketch the case of a planar configuration space C. See [16] for full details. We take C to be a rectangle in the plane containing polygonal obstacles B; the free space F is the closure of the set $C \setminus B$. (For simplicity we also assume that the obstacles do not meet the boundary of C.)

We first define the medial axis of a closed polygonal region P in the plane, 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 $\rho_P(x) = \mathrm{dist}(x, \partial P)$ is the distance to the boundary.[3] The *medial axis* $\mathrm{MA}(P)$ of P is defined to be the set of all points x of P whose associated $B_P(x)$ is maximal with respect to containment; i.e.,

$$\mathrm{MA}(F) = \{x \in F \mid \nexists\, y \in F \text{ with } B_F(x) \subsetneq B_F(y)\}.$$

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

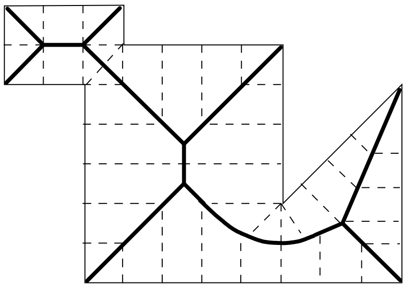

Figure 1: Medial axis of a polygon

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

It is easy to see that any multiple point will be in the medial axis. To see that the interior P° of a polygon may be continuously deformed or *retracted* onto the medial axis, note that any point not on the medial axis must be a simple point, i.e., it has a unique nearest point on the boundary ∂P. Such points x may be moved away from their nearest boundary points along the line connecting the two until the medial axis is reached; this defines the *canonical retraction map* $r_P : P^\circ \to \mathrm{MA}(P)$. See Figure 2.

Observe that it is possible to compute the image $r_P(x)$ of a single point x without actually computing the medial

[3] We define $\mathrm{dist}(x, S) = \inf_{y \in S} d(x, y)$ and $\mathrm{dist}(R, S) = \inf_{x \in R} \mathrm{dist}(x, S)$ where d denotes the euclidean distance.

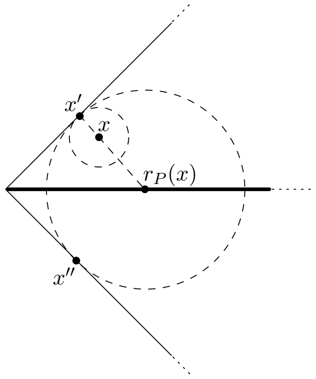

Figure 2: Image of a point x under the canonical retraction map.

axis first. We simply move the point x away from its nearest boundary point $x' \in \partial F$ until x has an additional nearest boundary point $x'' \in \partial F$. At that moment, the circle $B_F(x)$ contains the two boundary points x' and x'', so x has reached the medial axis. This can be carried out to compute the image to arbitrary precision using a bisection method; the only required geometric primitive is the ability to compute the nearest boundary point.

Returning to the motion planning example, we can now easily retract any sampled configuration in the free space F onto $\mathrm{MA}(F)$. See Figure 3(a).

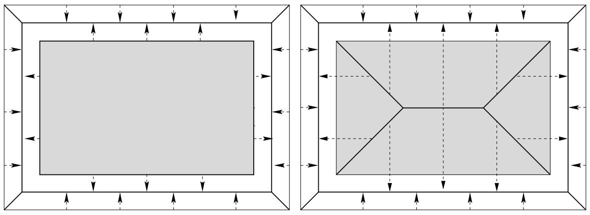

(a) (b)

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

From Figure 3(b), we can see how this map may be extended into the obstacle B as well. Any configuration in $B \setminus \mathrm{MA}(B)$ will have a unique nearest point on the boundary $\partial B \subseteq \partial F$. We can map such points through their nearest boundary points and into the free space F and then retract them to the medial axis of the free space. This map taking all of $C \setminus \mathrm{MA}(B)$ onto $\mathrm{MA}(F)$ is called the *extended retraction map*.

The MAPRM algorithm is essentially the following modification to uniform random sampling: rather than keeping only sampled free configurations, we keep *all* sampled configurations and retract them onto $\mathrm{MA}(F)$. We will show in Section 5 that this will increase the proportion of sampling in narrow passages.

Figure 4 shows an example of a free space containing a narrow corridor. Part (a) shows the result of sampling uniformly from the square until 100 nodes were obtained in the free space: this required generating 168 random configurations. Part (b) shows the result of sampling and retracting 100 nodes using the MAPRM algorithm. Most of the nodes in the corridor are nodes that were initially in collision and were retracted onto the medial axis. For this reason pushing the two obstacles closer together would not greatly affect the number of nodes MAPRM generates in the corridor.

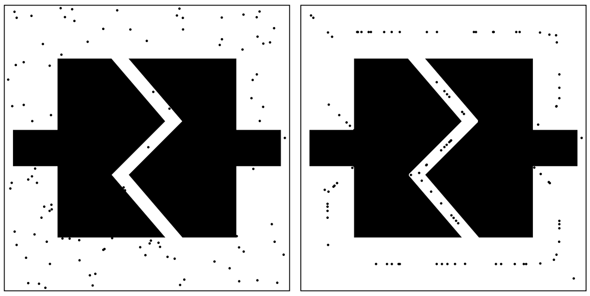

(a) (b)

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

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 the configuration space SE(3), which is six dimensional. Specifically, we want to retract all configurations, free or otherwise, onto the medial axis of the free space. Again, we would like to avoid explicit computation of the medial axis, but we also seek to avoid the costly computation of the C-obstacle boundaries in configuration space. Under a large class of metrics on the configuration space, the medial axis and retraction maps have a simple interpretation in terms of geometry in the workspace, allowing this calculation to be carried out without computation of the C-obstacle boundaries. For computational feasibility, our algorithm retracts all free configurations, but only a subset of the collision configurations. This subset is still sufficient to increase sampling in narrow passages.

4.1 Geometry on SE(3)

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 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, which imposes the necessary geometric structure, but there are still many possible choices.

A natural means of choosing a Riemannian metric is to require that it be independent of choice of coordinate system on the workspace and the workpiece, or equivalently, that it be invariant under changes in these coordinate frames. However, it has been shown that there is no Riemannian metric on SE(3) which is simultaneously invariant under changes of body frame and changes of world frame, see [11]. A metric invariant under changes in *either* the body frame or the world frame is easy to produce, however, see [15].

Rather than restricting our attention to a particular Riemannian metric, we will enumerate our assumptions about the metric being used and give a retraction algorithm which will always work under these conditions. Under such a metric, the medial axis and retraction maps will have a simple interpretation in terms of geometry and motions in the workspace. First we will introduce the notion of the medial axis for a subset of a complete Riemannian manifold.

4.2 The medial axis on a complete Riemannian manifold

We will not need to deal with many details of Riemannian geometry here, but a few facts are needed, see [4] for details. Roughly speaking, a *smooth manifold* M is a higher dimensional analog of a surface and a *Riemannian metric* defines an inner product on tangent vectors to M at a point; this gives us the ability to measure the length of a vector and the angle between a pair of vectors. In particular, integrating the length of a tangent vector to a smooth curve along the entire curve gives a measure of the length of the curve.

- A Riemannian metric induces a distance metric $d(\cdot,\cdot)$ in the metric space sense, i.e., a symmetric non-degenerate function satisfying the triangle inequality. The metric is obtained by assigning to each pair of points the infimum of the lengths of all piecewise smooth curves connecting the points.
- There is a distinguished set of smooth curves in M called *geodesics* which minimize the distance between any two nearby points on the path, along the entire length of the path.[4] The geodesics play roughly the same role played by lines in euclidean geometry.
- A *minimizing geodesic* between points p and q is a geodesic connecting p to q whose length is exactly $d(p,q)$.
- A Riemannian manifold is *complete* if closed and bounded sets (in the metric) are compact.[5] A complete Riemannian manifold has a minimizing geodesic connecting any given pair of points.

The plane with its usual euclidean geometry is an example of a complete Riemannian manifold: the geodesics are lines and there is always a line connecting any pair of points. Note that the plane with a single point removed is no longer complete: most of the interesting properties of the medial axis fail to be true in such a setting. The manifold we are ultimately interested in is SE(3) which will be a complete Riemannian manifold with our choice of Riemannian metric.

Let M be a complete Riemannian manifold. Due to the induced metric space structure on M, we can define the medial axis as before: given a proper closed subset F of M, we can again define the boundary distance function $\rho_F(x) = \mathrm{dist}(x, \partial F)$, and maximal ball at a point $B_F(x) = \overline{B}(x, \rho_F(x))$. We define the *medial axis* as before:

$$\mathrm{MA}(F) = \{x \in F \mid \nexists\, y \in F \text{ with } B_F(x) \subsetneq B_F(y)\}.$$

Our retraction algorithm will work essentially as before. (See Figure 2.) Let $x \in F$, x not in $\mathrm{MA}(F)$, and suppose that there is a unique point $r_F(x) \in \mathrm{MA}(F)$ such that $B_F(x) \subset B_F(r_F(x))$. Then x will be on a minimizing geodesic connecting $r_F(x)$ to the unique nearest boundary point x' of x, i.e., there is a unit speed ($|\gamma'| \equiv 1$) geodesic $\gamma : \mathbb{R} \to F$ with $\gamma(0) = x'$, $\gamma(d(x', x)) = x$, and $\gamma(d(x', r_F(x))) = r_F(x)$. It is clear that x' is not a nearest boundary point of $\gamma(t)$ when $t > d(x', r_F(x))$, for otherwise $B_F(\gamma(t))$ would contain $B_F(r_F(x))$. This allows us to again use a bisection method to locate medial axis points to arbitrary precision, provided we can find the closest point on ∂F to a given point.

[4] Geodesics do not necessarily minimize the distance between any two points on them, only between points that are sufficiently close together.

[5] Completeness is usually defined in terms of the exponential map: this is equivalent via the Hopf-Rinow Theorem, see [4].

As in the planar case, we can extend this map to almost the entire manifold M by mapping points outside F through their nearest boundary point and slightly into the interior of F; the retraction algorithm described above can then be applied. If we let $B = \overline{M \setminus F}$, any point in $B \setminus \mathrm{MA}(B)$ will have a unique minimal geodesic to the boundary which we can follow to push the point through the boundary $\partial F = \partial B$ and slightly into the interior of F. We again call this the *extended retraction map.*

Note that in order to carry out this procedure we require the ability to compute the nearest boundary point of a given point, as well as the ability to compute and follow a minimizing geodesic connecting them.

4.3 The medial axis on SE(3)

Returning to SE(3), there are many possible choices of Riemannian metric, however some choices seem bizarre from the point of view of motion planning. Rather than restrict to a particular metric, we make two assumptions about the metric and show how the extended retraction map may be computed in such a case.

Given a workpiece U and obstacle V in $\mathbb{R}^3$, we assume we are provided with a Riemannian metric on SE(3) satisfying the following two properties:

1. Translation is always a minimizing geodesic; i.e., given two elements $(R, p_1), (R, p_2) \in C$, a shortest path between them is given by
$$\gamma(t) = (R, (1-t)p_1 + tp_2).$$
We assume the length of γ in this case is the usual euclidean length $|p_2 - p_1|$.

2. Take U to be at a particular free configuration $g \in F$. We assume that the shortest path from g to ∂F (the set of contact configurations) is a pure translation (i.e., with no rotation) toward the nearest workspace obstacle point.

The second property essentially amounts to the assumption that rotations are weighted heavily enough that among all transformations taking a particular fixed point of U to a prescribed location in the workspace, pure translation is always the cheapest. Note that this assumption depends on the particular object U.

As an example of a metric satisfying these properties, we consider the *product metric* on SE(3). Note that $\mathbb{R}^3$ and SO(3) each come equipped with a Riemannian metric: $\mathbb{R}^3$ has the usual euclidean metric, and SO(3) has a well-known bi-invariant metric which is unique up to scale; we assume this has been normalized, so that, e.g., a rotation about an axis by $\theta \in [0, 2\pi]$ radians has distance θ from the identity. There is a family of similar product metrics on the product manifold SE(3); these will be world-frame (left) invariant under multiplication in SE(3). Such a product metric is fixed once we decide the weights the translation and rotation components should receive. Our first assumption forces the translation component to have weight 1. For a particular workpiece U, which is assumed to be a bounded set, we weight the rotation component by $R > 0$ at least as large as the distance from the (body frame) origin to the most distant point of U. See [15] for details on the product metric on SE(3).

4.4 The Algorithm

We now assume that we have $C = \mathrm{SE}(3)$ with a Riemannian metric satisfying the above properties. Our assumptions on the metric have the following consequences for the medial axis:

- Let $g \in F$. A nearest contact configuration to g (i.e., a nearest point of ∂F) is obtained by first finding a closest pair of points $p \in g \cdot U$ and $q \in V$. A nearest contact configuration g' is given by translating p to q; i.e., $g' = (I, q-p) \cdot g$.
- Let $g \in F$. If $g \cdot U$ has two distinct nearest obstacle points, i.e., there are two distinct points $y_1, y_2 \in V$ such that $\mathrm{dist}(V, g\cdot U) = \mathrm{dist}(y_1, g\cdot U) = \mathrm{dist}(y_2, g\cdot U)$, then g is on the medial axis of F.

We assume U and V are given as unions of polyhedra. The resulting algorithms are given in Algorithms 4.1 and 4.2. Our brute-force algorithm for computing the nearest contact configuration for a collision configuration is essentially an exhaustive search for the shortest translation.[6] We ignore the possibility (a probability 0 event) that a sampled configuration actually lands on $\mathrm{MA}(B)$.

Algorithm 4.2 finds the nearest contact configuration g' to a given configuration g provided g' puts the workpiece U into contact with the obstacle V at only a single point.[7] If the nearest contact configuration g' puts U into contact with V at more than one point, Algorithm 4.2 does not return g' but instead the returns the shortest *translation* from g to a boundary configuration (if possible). This is done for the sake of computational feasibility. A configuration g' placing U in contact with V at more than one point is already on the medial axis of F, so this shortcut only modifies the sampling at contact configurations on the medial axis of F, which we expect to be a small subset of $\mathrm{MA}(F)$.

5 Analysis

In this section we state in precise terms why the sampling rate is increased in narrow passages, and we observe that the amount of increase in the sampling rate depends in the "thickness" of the obstacles bounding the corridor.

Overall performance results, including the influence of the choice of metric on the performance, have not been developed.

5.1 Sampling in narrow corridors

It is fairly clear that MAPRM will increase the sampling rate in small corridors: a nonzero volume of nearby configurations in B will be pushed into the corridor, increasing the sampling by a constant that is essentially independent of the volume of the corridor. In this section we make this idea somewhat more precise.

The medial axis provides a convenient definition for what is meant by a corridor. Let $r_F : F^\circ \to \mathrm{MA}(F)$ be the canonical retraction map. A *corridor* in F is a connected open subset S of F° that contains the retracted images of all of its points (i.e., $r_F(S) \subseteq S$) while also including any points

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

[7] One can show that if g' is a nearest contact configuration to g and g' puts U into contact with V at only one point, then the minimizing geodesic connecting g to g' must be a pure translation (assuming polyhedral U and V).

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.

repeat
  repeat
    Sample a configuration (R, p) from C.
    Run Algorithm 4.2 which returns either the nearest contact configuration (R, q) to (R, p), or **failure**.
  until Algorithm 4.2 succeeds.
  if (R, p) is free **then**
    Take the retraction direction $\vec{v}$ to be $\overrightarrow{qp}$, and let the start point s be p.
  else
    Take the retraction direction $\vec{v}$ to be $\overrightarrow{pq}$, and let the start point s be q.
  end if
  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 free space's medial axis.
until N vertices have been output
For each pair of vertices: if the pair can be connected with the local planner, insert an edge into the graph connecting them.

Algorithm 4.2 Finding 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**.

if (R, p) is a free configuration **then**
  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.
else
  For each feature (vertex, edge, face) of U, and each feature of the V, find the configuration (R, q) with smallest $|p - q|$ that puts these features in contact in a single point.
  Such a configuration may put U strictly in collision (not just contact) with an obstacle; discard any such configurations.
  If no configurations remain, output **failure**.
  Otherwise, output a remaining configuration (R, q) with smallest $|p - q|$.
end if

that will be retracted into S under r_F (i.e., $r_F^{-1}(S) \subseteq S$). These conditions essentially ensure that S is bounded on "all" sides by obstacles.

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 the closure $\overline{S}$ will be mapped into $\overline{S}$. Let $b_B : B \setminus \mathrm{MA}(B) \to \partial B$ be the map that takes each point to its nearest boundary point. (This is well-defined since any multiple point will be in $\mathrm{MA}(B)$.) The volume of points that map into the $\overline{S}$ under the full retraction map is:

$$\mu(S) + \mu(b_B^{-1}(\partial S)),$$

where μ is a volume measure[8] on C.

Using uniform random sampling from C (with respect to μ) with collisions discarded, the probability that a single sampled configuration is contained in S is:

$$\frac{\mu(S)}{\mu(C)}. \tag{1}$$

Assuming the medial axis $\mathrm{MA}(B)$ has measure zero, the probability that a single random sample from C is retracted by MAPRM into $\overline{S}$ is:

$$\frac{\mu(S) + \mu(b_B^{-1}(\partial S))}{\mu(C)} \tag{2}$$

Assuming the term $\mu(b_B^{-1}(\partial S))$ is nonzero, (2) will exceed (1); this shows that the sampling rate in the corridor is increased, see Figure 5. Note that the term $\mu(b_B^{-1}(\partial S))$ only involves the boundary of the corridor S and does not involve the volume of S.

[8] A volume measure is provided, for example, by the Riemannian metric on SE(3). We assume F, B, and S, etc., are μ-measureable.

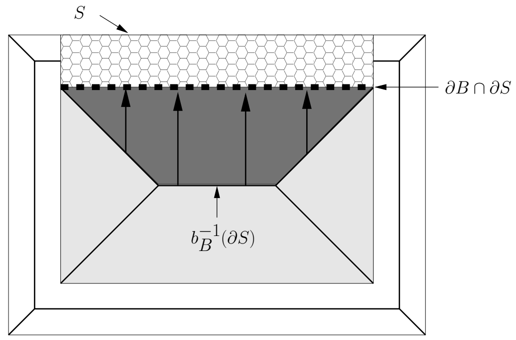

Figure 5: The hatched area is S, the heavy dashed line is $\partial B \cap \partial S$, and the dark shaded area is $b_B^{-1}(\partial S)$.

Intuitively, we may regard the function $A \mapsto \mu(b_B^{-1}(A))$ defined on subsets A of the boundary of B as a measure of how "thick" the obstacle is behind A. Equation (2) shows that our algorithm gives a greater increase in the sampling rate for corridors that happen to be bounded by thicker obstacles.

In practice, taking a single uniform random sample from C and running a collision check is much faster than a random sample followed by retraction using Algorithm 4.1. However, in our preliminary experiments (Section 6), MAPRM generates *free* nodes at a much higher rate, giving faster overall performance.

The above analysis requires some rigorous justification, particularly the assertion that $\mu(b_B^{-1}(\partial S))$ is nonzero. It is easy to see that this will be true, e.g., for moving a point among polygons in the plane, but an argument for more general B is required. Also, as discussed in Section 4.4, Algorithm 4.2 is only guaranteed to compute the nearest contact configuration if this configuration is not on the medial axis of F, so the estimate $\mu(b_B^{-1}(\partial S))$ should be modified to be $\mu(b_B^{-1}(\partial S \setminus \mathrm{MA}(F)))$. This *can* make a significant difference, however we still expect the term $\mu(b_B^{-1}(\partial S \setminus \mathrm{MA}(F)))$ to be nonzero. Full justification will be discussed in a future paper.

5.2 Complexity

The brute-force approach to find the nearest contact configuration in Algorithm 4.2 consists of enumerating all possible single-point contacts between features of obstacle V and the workpiece U in the given orientation, sorting them according to translation distance from the given configuration, and finding a free one with minimal translation distance. This has time complexity $O(n_U n_V \log(n_U n_V) + n_U n_V t_{cd}(n_U, n_V))$ and uses $O(n_U n_V)$ storage, where n_U and n_V are the number of features of U and V, and $t_{cd}(n, m)$ is the collision detection time for two objects of size n and m, respectively.

If both the workpiece U and obstacle V are convex, the shortest translation that renders the interiors of U and V disjoint can be found in $O(n_U n_V)$ time [5].

6 Experimental results

In this section we present experimental results to show the sampling increases obtained by MAPRM on two examples requiring traversal of narrow corridors.

We implemented MAPRM for rigid bodies using V-Clip [12] to provide collision detection and closest pair calculations. We used a single local planner: translation with simultaneous rotation about the principal axis of rotation in the body frame. In using probabilistic roadmap methods in practice, the the time to connect nodes usually overwhelms the time spent in the sampling phase. As a result, connections are not usually attempted between all pairs of nodes, but only between promising pairs of nodes according to some heuristic. In our examples, so few nodes are generated that it was feasible to attempt connections between all pairs necessary to determine the components of the roadmap.

6.1 Narrow corridor example

Our first example is shown in Figure 6. The workpiece is a cube of side length 2; the obstacle is a solid cube of side 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. The retraction and collision checking calculations were carried out with a spatial tolerance of 0.01. 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.[9] 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.

[9] Experiments were run on a MIPS R10000 processor running at 200 Mhz.

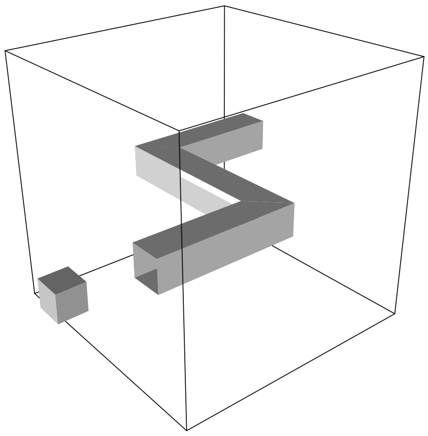

Figure 6: Rigid body example. The workpiece is the small cube and the obstacle is a solid block with the indicated corridor cut through it.

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	1.706	22.652

6.2 Wider corridor example

We also compared performance on a simpler problem: the same obstacle block but with a workpiece cube of side length 1.5 units (rather than 2 units). Mean results for 15 trials are given in Table 2. In this case both methods solved the problem very quickly, but uniform sampling was quicker on average than MAPRM. Still, MAPRM generated free nodes at a higher rate than uniform sampling. Again, note the significantly larger number of roadmap nodes required by MAPRM.

We should also note that the variance from the mean was very high in both of these examples. For example, in the narrow corridor case, the mean sampling time for MAPRM was 690 seconds with standard deviation 163; the mean sampling time with uniform sampling was 7875 seconds with standard deviation 2514. This is probably due to the fact that a few "key" nodes near the corners are sufficient to solve the problem. In more realistic examples, connections are only attempted between nearby pairs of nodes, so we would expect these deviations to decrease.

Table 2: Experimental results for smaller workpiece

	MAPRM	Uniform
Sampling time	14.6	10.5
Connection time	12.4	4.4
Total preprocessing time	27.0	14.9
# roadmap nodes required	109	57
# random config. sampled	881	157,333
Mean time to generate a node	0.133	0.184

7 Conclusions

We have described a new sampling method for probabilistic roadmap motion planning for a rigid body. The method retracts sampled configurations onto the medial axis of the free space while avoiding explicit calculation of the C-obstacles and the medial axis. A theoretical explanation was given to show the this retraction method increased the probability that a sampled (and retracted) node lands in a corridor. Experimental results were given to show that this does increase the sampling rate in simple cases.

Although the method does increase the sampling rate, it employs more complicated geometric calculations than the standard uniform sampling and is consequently more difficult to implement. It also takes longer to run per sampled node, but in our examples the rate of generation of free nodes was actually higher. For larger problems, we expect the time required for the calculation of the nearest contact configuration to become even more significant. For less crowded environments, uniform random sampling is likely to perform better, suggesting that in general some hybrid algorithm may outperform either algorithm individually.

It is unclear how this method would extend to articulated robots, an area in which probabilistic methods have been very successful. Finally, more theoretical analysis is required to understand the overall effectiveness or probability of success of the method and to understand the effects of the choice of configuration space metric.

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] M. P. a. do Carmo. *Riemannian geometry.* Birkhäuser Boston Inc., Boston, MA, 1992. Translated from the second Portuguese edition by Francis Flaherty.

[5] D. Dobkin, J. Hershberger, D. Kirkpatrick, and S. Suri. Computing the intersection-depth of polyhedra. *Algorithmica*, 9:518–533, 1993.

[6] 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.

[7] L. Kavraki. *Random Networks in Configuration Space for Fast Path Planning.* PhD thesis, Stanford Univ., Stanford, CA, 1995.

[8] L. E. Kavraki, J.-C. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot path planning. In *Proc. 27th Annu. ACM Sympos. Theory Comput.*, pages 353–362, 1995.

[9] 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.

[10] J.-C. Latombe. *Robot Motion Planning.* Kluwer Academic Publishers, Boston, 1991.

[11] J. Loncaric. *Geometrical Analysis of Compliant Mechanisms in Robotics.* PhD thesis, Harvard Univ., Cambridge, MA, 1985.

[12] B. Mirtich. V-Clip: Fast and robust polyhedral collision detection. Technical Report TR97-05, MERL, 201 Broadway, Cambridge, MA 02139, USA, July 1997.

[13] R. M. Murray, Z. Li, and S. S. Sastry. *A Mathematical Introduction to Robotic Manipulation.* CRC Press, Boca Raton, FL, 1994.

[14] M. Overmars and P. Svestka. A probabilistic learning approach to motion planning. In *Proc. Workshop on Algorithmic Foundations of Robotics*, pages 19–37, 1994.

[15] 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.

[16] 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.