

Probabilistically Complete Planning with End-Effector Pose Constraints

Dmitry Berenson[†]

[†]The Robotics Institute

Carnegie Mellon University

5000 Forbes Ave., Pittsburgh, PA, 15213, USA

dberenso@cs.cmu.edu

Siddhartha S. Srinivasa^{‡†}

[‡]Intel Labs Pittsburgh

4720 Forbes Ave., Suite 410

Pittsburgh, PA, 15213, USA

siddhartha.srinivasa@intel.com

Abstract—We present a proof for the probabilistic completeness of RRT-based algorithms when planning with constraints on end-effector pose. Pose constraints can induce lower-dimensional constraint manifolds in the configuration space of the robot, making rejection sampling techniques infeasible. RRT-based algorithms can overcome this problem by using the *sample-project* method: sampling coupled with a projection operator to move configuration space samples onto the constraint manifold. Until now it was not known whether the sample-project method produces adequate coverage of the constraint manifold to guarantee probabilistic completeness. The proof presented in this paper guarantees probabilistic completeness for a class of RRT-based algorithms given an appropriate projection operator. This proof is valid for constraint manifolds of any fixed dimensionality.

I. INTRODUCTION

This paper focuses on the completeness properties of RRT-based algorithms that plan paths for tasks involving pose constraints on a manipulator’s end-effector. The pose of a manipulator’s end-effector is represented by a point in $SE(3)$, the six-dimensional space of rigid spatial transformations. Many practical manipulation tasks, such as carrying a pitcher of liquid or sliding a heavy object on a table, impose constraints on the motion of a robot’s end-effector.

Our framework for planning with such constraints comprises of two parts: a representation of pose constraints which we term Task Space Regions (TSRs), and a sampling-based planner, the Constrained BiDirectional RRT (CBiRRT) [1] that plans paths on constraint manifolds. Although we have demonstrated that the TSR representation can describe many useful tasks and the CBiRRT is capable of planning for these tasks, we have yet to demonstrate that it is probabilistically complete. A path planner is called *probabilistically complete* if, given a problem that is solvable, the probability that the planner solves the problem goes to 1 as the running time goes to infinity. To our knowledge, probabilistic completeness has yet to be shown for *any* sampling-based planner which plans with pose constraints. The central difficulty lies in the coverage of the constraint manifold. Depending on the definition of an end-effector pose constraint, it can induce a variety of manifolds in the robot’s configuration space (C-space). If these manifolds have non-zero volume in the C-space (see Figure 1f) it is straightforward to show that an RRT-based algorithm is probabilistically complete because rejection sampling in the C-space will eventually place samples inside of the manifold.

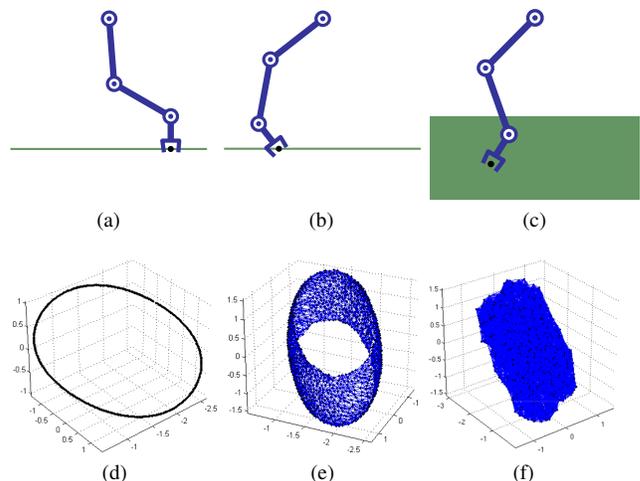


Fig. 1. Three pose constraints and their corresponding C-space constraint manifolds for a 3-link manipulator. (a) The end-effector must be on the line pointing downward. (b) The end-effector must be on the line with an orientation $\pm 0.7\text{rad}$ of downward. (c) The end-effector must be in the green rectangle with an orientation $\pm 0.7\text{rad}$ of downward. (d-f) show graphs created from sampling on the manifolds corresponding to the constraints in (a-c), respectively. Black points are nodes and blue lines are edges. The manifold in (f) has non-zero volume in the C-space, the manifolds in (d) and (e) do not.

However, if a pose constraint induces a lower-dimensional manifold, i.e. one that has zero volume in the C-space (see Figures 1d and 1e), rejection sampling in the C-space will not generate a sample on the constraint manifold. One way to overcome this difficulty is to use the *sample-project* method. This method consists of two steps: first we take a sample in the C-space and then we project that sample to the constraint manifold using a projection operator. Until now the properties of the projection operator sufficient to achieve constraint manifold coverage have not been known. Thus the probabilistic completeness of RRT-based algorithms which use the sample-project method has yet to be shown.

Our proof for probabilistic completeness has two parts: first, we present a set of properties for the projection operator and prove that these properties allow the sample-project method to cover the constraint manifold. Second, we describe a class of RRT-based algorithms (such as CBiRRT [1] and TCRRT [2]), which use such a projection operator and prove that they are probabilistically complete. We conclude with a discussion of the implications of the proof for projection operators and mixed-dimensional constraint manifolds.

II. BACKGROUND

Our method of planning with pose constraints builds on several developments in motion planning and control research in robotics. The planning algorithms we describe are based on the Rapidly-exploring Random Tree (RRT) algorithm by LaValle and Kuffner [3]. For meeting constraints on end-effector pose, this paper considers projection methods similar to those used for inverse kinematics. Iterative inverse kinematics algorithms use projection methods based on the pseudo-inverse or transpose of the Jacobian to move a robot's end-effector closer to some desired pose (e.g. [4]).

An alternative projection operator is Randomized Gradient Descent (RGD) [5], which uses random sampling of the C-space to iteratively project a sample toward a closed-chain constraint. Stilman [2] showed that when RGD is extended to work with more general pose constraints it is significantly less efficient than Jacobian pseudo-inverse projection and it is sometimes unable to meet more stringent constraints. Thus we focus only on Jacobian-based methods in our proof.

Our proof only considers algorithms that plan in the C-space, unlike the task-space planners in [6][7][8], which assign a single configuration to each task-space point (from a potentially infinite number of possible configurations). Though these methods are often efficient, they are not probabilistically complete.

Finally, our proof of manifold coverage hinges on the concept of *self-motion manifolds*, which were described by Burdick [9]. A self-motion manifold is the set of configurations in C-space which place the end-effector of the robot in a certain pose.

III. GENERAL PROPERTIES

Before we begin the proof for probabilistic completeness, we discuss some general properties of manifolds that will be useful in the proof. For brevity, we will only state parts of definitions that are relevant to our proof.

A topological *manifold* is a second countable Hausdorff space where every point has a neighborhood homeomorphic to an open Euclidean n -ball, where n is allowed to vary. Manifolds can be disjoint, with each piece of the manifold called a *connected component*. Manifolds that have a fixed n (i.e. n -dimensional manifolds) are called *pure* manifolds.

Lemma 3.1: Consider manifolds \mathcal{A} and \mathcal{B} . $(\mathcal{A} \cap \mathcal{B})$ is n -dimensional if the following conditions are true:

- 1) \mathcal{A} and \mathcal{B} are both n -dimensional
- 2) \mathcal{A} and \mathcal{B} are both submanifolds of the same n -dimensional manifold.
- 3) $(\mathcal{A} \cap \mathcal{B}) \neq \emptyset$

Proof: Consider a point $p \in (\mathcal{A} \cap \mathcal{B})$. \mathcal{A} and \mathcal{B} both contain the same n -dimensional ball centered at p by the first and second conditions. $(\mathcal{A} \cap \mathcal{B})$ is the union of all such n -dimensional balls for all p . This union is n -dimensional. ■

Let μ_n be a measure of volume in an n -dimensional space. If the volume of a manifold in an embedding space is zero, then the probability of generating a sample on the manifold by rejection sampling in the embedding space is zero. Conversely if the volume of a manifold in an embedding space

Name	Symbol	Dimension
C-space	\mathcal{Q}	n
Configuration	q	0
Constraint Manifold	\mathcal{M}	$m = n - (r - d)$
Reachability	\mathcal{R}	r
Pose	$x(q)$	0
Task Constraint	\mathcal{T}	$d \leq r$

TABLE I

DEFINITIONS USED THROUGHOUT THE PROOF.

is *not* zero, then the probability of generating a sample on the manifold by rejection sampling in the embedding space will go to 1 as the number of samples goes to infinity.

Definition A sampling method *covers* a manifold if it generates a set of samples such that any open n -dimensional ball contained in the manifold contains at least one sample.

IV. DEFINITIONS

Let the reachable manifold of end-effector poses of the given manipulator be $\mathcal{R} \subseteq SE(3)$. \mathcal{R} is defined by the Forward Kinematics function of the robot:

$$x : \mathcal{Q} \rightarrow \mathcal{R} \quad (1)$$

where \mathcal{Q} is the C-space. x is always surjective and can be one-to-one or many-to-one, depending on the manipulator.

We restrict our proof to manipulators whose \mathcal{Q} and \mathcal{R} are both pure manifolds. This restriction is necessary to meet the conditions of Lemma 3.1, yet allows many manipulators commonly used today, such as serial-chain manipulators and humanoids. In this paper, we also restrict our focus to manipulators with no non-holonomic constraints, though we hypothesize that our proof can extend to the non-holonomic case as well.

Let \mathcal{Q} be n -dimensional, where n is the number of DOF of the manipulator. We will parameterize $SE(3)$ locally using three variables for translation and three for rotation, i.e. a pose will be a vector in \mathbb{R}^6 . Let \mathcal{R} be r -dimensional, where $r \leq 6$.

Let the manifold of end-effector poses allowable by the task be $\mathcal{T} \subseteq \mathcal{R}$. Let this manifold have dimensionality $d \leq r$. We will assume that d is fixed, though we will discuss the implications of allowing d to vary in Section VII. Let the manifold of configurations that place the robot's end-effector in \mathcal{T} be $\mathcal{M} \subseteq \mathcal{Q}$. \mathcal{M} is the union of all self-motion manifolds that map to a pose in \mathcal{T} :

$$\mathcal{M} = \bigcup_{t \in \mathcal{T}} \{q \in \mathcal{Q} \mid x(q) = t\} \quad (2)$$

\mathcal{M} has dimensionality $m = n - (r - d)$.

We will be using the (weighted-)Euclidian distance metric on $SE(3)$, which we will denote as *dist*. This distance metric has the property that each pose in \mathcal{T} has at least an $(r - d)$ -dimensional *Voronoi cell* in \mathcal{R} . A Voronoi cell is the set of points that are closer to a certain point than to any other given a distance metric [10].

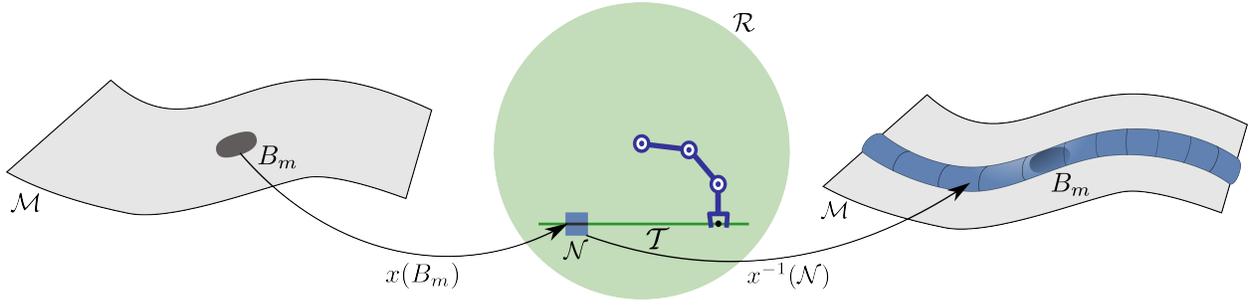


Fig. 2. An example showing the process used to define $x^{-1}(\mathcal{N})$. $d = 1$, $r = 2$, and $n = 3$.

Table I summarizes the definitions and dimensionalities of manifolds used in the proof.

V. PROOF OF MANIFOLD COVERAGE BY THE SAMPLE-PROJECT METHOD

When the manifold of allowable configurations in the C-space is of the same dimensionality as the C-space, samples on the manifold can be generated by *rejection sampling* in the C-space. Rejection sampling is simply the process of generating a sample in the C-space and then checking if it is an allowable configuration. However, when \mathcal{M} is of a lower dimension than the C-space, the probability of generating a sample on the manifold through rejection sampling is 0.

The *sample-project* method is an approach which can generate samples on manifolds of a lower dimension. This method first produces a sample in the C-space and then projects that sample onto \mathcal{M} using a *projection operator* $P : \mathcal{Q} \rightarrow \mathcal{M}$. In order to show that an algorithm using the sample-project method is probabilistically complete, we must first show that this method covers \mathcal{M} .

This section will show that the sample-project method does indeed cover \mathcal{M} if the projection operator has the following properties:

- 1) $P(q) = q$ if and only if $x(q) \in \mathcal{T}$
- 2) If $x(q_1)$ is closer to $x(q_2) \in \mathcal{T}$ than to any other point in \mathcal{T} and $\text{dist}(x(q_1), x(q_2)) < \epsilon$ for an infinitesimal $\epsilon > 0$, then $x(P(q_1)) = x(q_2)$.

The first property guarantees that a configuration that is already in \mathcal{M} will project to itself. The second property ensures that any pose in \mathcal{T} can be chosen for projection. We will describe the underlying mechanics of the projection operator in Section V-B.

If $d = r$ then $\mu_n(\mathcal{M}) > 0$. By the first property of P , a sample placed in \mathcal{M} will project to itself. Thus the sample-project method will cover \mathcal{M} by the same principle as rejection sampling.

The remainder of this section will focus on proving coverage when $d < r$, i.e. when $\mu_n(\mathcal{M}) = 0$. Consider an open m -dimensional ball $B_m(q) \subseteq \mathcal{M}$ for any $q \in \mathcal{M}$. Note that *open* in this context refers to the openness of the set with respect to \mathcal{M} . For notational simplicity, let B_m represent any $B_m(q)$ for any such q . We will show that the sample-project method places a sample in any B_m as the number of iterations goes to infinity, thus covering of \mathcal{M} .

Consider an n -dimensional manifold $\mathcal{C}(B_m) = \{q : P(q) \in B_m, q \in \mathcal{Q}\}$. If such a \mathcal{C} exists for any B_m , the sample-project method will place a sample inside \mathcal{C} with probability greater than 0 (because \mathcal{C} is n -dimensional) and that sample will project into B_m . This will guarantee coverage of \mathcal{M} as the number of iterations goes to infinity. But how do we guarantee that such a \mathcal{C} exists for any B_m ?

We will show that \mathcal{C} can be defined as the intersection of two n -dimensional manifolds, $x^{-1}(\mathcal{N})$ and $\mathcal{U}\mathcal{H}$, both of which must exist for any B_m (notation will be explained in subsequent subsections). The following subsections describe each of these manifolds and show that $(x^{-1}(\mathcal{N}) \cap \mathcal{U}\mathcal{H})$ must be n -dimensional and all configurations in $(x^{-1}(\mathcal{N}) \cap \mathcal{U}\mathcal{H})$ must project into B_m . Thus $(x^{-1}(\mathcal{N}) \cap \mathcal{U}\mathcal{H})$ meets the requirements of \mathcal{C} , which completes the proof of coverage.

A. To Task Space and Back Again: Defining $x^{-1}(\mathcal{N})$

In this subsection we will define a manifold $x^{-1}(\mathcal{N})$, which projects into a set of self-motion manifolds $\mathcal{S} \subseteq \mathcal{M}$ that intersects B_m . We will do this by mapping B_m into task space, constructing a manifold of poses that project into $x(B_m)$ and then mapping that manifold back into C-space (see Figure 2).

Mapping B_m into task space yields a d -dimensional manifold $x(B_m) \subseteq \mathcal{T}$. Let us define a task-space manifold $\mathcal{N}(x(B_m)) \subseteq \mathcal{R}$. \mathcal{N} is constructed in two steps. First, take the union of a set of r -dimensional open balls centered at every $x(q) \in x(B_m)$. Second, remove all poses that are closer to a $x(q) \in (\mathcal{T} - x(B_m))$ than to any $x(q) \in x(B_m)$ and remove all equidistant poses. Formally, \mathcal{N} is defined as:

$$\begin{aligned} \mathcal{U} &= \bigcup_{x(q) \in x(B_m)} B_r(x(q)) \\ \mathcal{N} &= \mathcal{U} - \{p \in \mathcal{U} \mid \exists p_1 \in (\mathcal{T} - x(B_m)) \\ &\quad \text{dist}(p, p_1) \leq \inf_{p_2 \in x(B_m)} \text{dist}(p, p_2)\} \end{aligned} \quad (3)$$

where $B_r(x(q))$ is an open r -dimensional ball centered at $x(q)$ with radius ϵ . \mathcal{N} has the following properties:

- 1) \mathcal{N} is r -dimensional.
- 2) \mathcal{N} contains $x(B_m)$.
- 3) $x(P(q)) \in x(B_m)$ for all $\{q \in \mathcal{Q} \mid x(q) \in \mathcal{N}\}$

The first property follows from the fact that the Voronoi cell of any $x(q) \in \mathcal{T}$ is $(r - d)$ -dimensional and the fact that

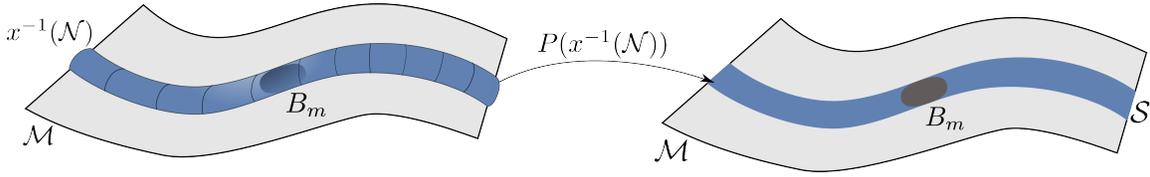


Fig. 3. An example of $P(x^{-1}(\mathcal{N})) = \mathcal{S}$. $d = 1$, $r = 2$, and $n = 3$.

$x(B_m)$ is d -dimensional. The second property is clear by construction. The third property is guaranteed by the second property of the projection operator.

Define the C-space manifold $x^{-1}(\mathcal{N}) = \{q \in \mathcal{Q} \mid x(q) \in \mathcal{N}\}$. Since \mathcal{N} is r -dimensional, x^{-1} will map it to an n -dimensional manifold of configurations in C-space (recall that x is surjective). $x^{-1}(\mathcal{N})$ has the following properties:

- 1) $x^{-1}(\mathcal{N})$ is n -dimensional.
- 2) $x^{-1}(\mathcal{N})$ contains B_m

Now the question becomes, where do configurations in $x^{-1}(\mathcal{N})$ project to? It would be convenient if $P(x^{-1}(\mathcal{N})) = B_m$, however this is not always the case. The reason is that a single pose can correspond to more than one configuration in C-space, i.e. to a self-motion manifold. This occurs when the manipulator is redundant, for instance.

Let us now define where $x^{-1}(\mathcal{N})$ projects to. By the third property of \mathcal{N} , we know that all poses in \mathcal{N} will project into $x(B_m)$. The analog of $x(B_m)$ in C-space will be a manifold $\mathcal{S} \subseteq \mathcal{M}$. \mathcal{S} is defined as the manifold of configurations which map into $x(B_m)$; i.e. $\mathcal{S} = \{q \in \mathcal{M} \mid x(q) \in x(B_m)\}$. Thus we can state a third property of $x^{-1}(\mathcal{N})$ (see Figure 3):

- 3) $P(x^{-1}(\mathcal{N})) = \mathcal{S}$.

We will use \mathcal{S} to show coverage of \mathcal{M} in Section V-C, but first we must show that \mathcal{S} contains all self-motion manifolds that intersect B_m and only those self-motion manifolds.

Lemma 5.1: \mathcal{S} has the following properties:

- 1) If a self-motion manifold intersects B_m , it is a subset of \mathcal{S} .
- 2) If a self-motion manifold does not intersect B_m , it is not a subset of \mathcal{S} .
- 3) $\mathcal{S} \neq \emptyset$.

Proof of 1st Property: Recall that, for any self-motion manifold $s \subseteq \mathcal{M}$, $x(s) = x(q)$ for any $q \in s$; i.e. all configurations on the self-motion manifold map to the same pose. If there exists $q \in (s \cap B_m)$, then $x(q) \in x(B_m)$, which entails $x(s) \in x(B_m)$. Thus by definition of \mathcal{S} , $s \subseteq \mathcal{S}$. ■

Proof of 2nd Property: We will show this by contradiction. Suppose there is a self-motion manifold $s \subseteq \mathcal{S}$ that does not intersect B_m . By definition of \mathcal{S} , $x(s)$ would have to be in $x(B_m)$. However, if $x(s) \in x(B_m)$, there is some configuration $q \in B_m$ such that $x(q) = x(s)$. This entails that $q \in s$ and thus $q \in (s \cap B_m)$, but this contradicts the assumption that s does not intersect B_m . ■

Proof of 3rd Property: \mathcal{M} is composed of all self-motion manifolds for all poses in \mathcal{T} . Since $\mu_m(B_m) > 0$, B_m must intersect at least one self-motion manifold and by the first property of \mathcal{S} , $\mathcal{S} \neq \emptyset$. ■

B. Projection onto a ball on a self-motion manifold

We have shown that $P(x^{-1}(\mathcal{N})) = \mathcal{S}$ and described some properties of \mathcal{S} , however we have not stated where on \mathcal{S} a projected configuration will go. It is possible that a $q \in x^{-1}(\mathcal{N})$ will project to a configuration outside of B_m because \mathcal{S} may contain configurations outside of B_m .

In order to show that the probability of placing a sample inside B_m using the sample-project method is greater than 0, we need to show that there exists an n -dimensional manifold around a ball on a self-motion manifold that projects into that ball. The purpose of this subsection is to prove this property of self-motion manifolds. The remainder of this subsection will consider a self-motion manifold in isolation in order to show this property.

Let us now look closer at the mechanism of projection used by P . In this paper, we focus on projection operators based on Jacobian pseudo-inverse or Jacobian transpose. These kinds of projection operators step towards a pose target and this process can be written as a differential equation:

$$\frac{dq}{dt} = f(q(t)), t \in [0, \infty) \quad (4)$$

f satisfies the Lipschitz condition:

$$\|f(q_1) - f(q_2)\| \leq L \|q_1 - q_2\| \quad (5)$$

where $\|\cdot\|$ denotes any p -norm and L is the Lipschitz constant [11].

An equilibrium point \bar{q} of Eqn.4 satisfies $f(\bar{q}) = 0$.

Definition The equilibrium point \bar{q} is

- *stable* if, for each $\epsilon > 0$, there is a $\delta = \delta(\epsilon) > 0$ such that

$$\|q(0) - \bar{q}\| < \delta \Rightarrow \|q(t) - \bar{q}\| < \epsilon, \forall t \geq 0$$

- *unstable* if it is not stable
- *asymptotically stable* if it is stable and δ can be chosen such that

$$\|q(0) - \bar{q}\| < \delta \Rightarrow \lim_{t \rightarrow \infty} q(t) = \bar{q}$$

Asymptotic stability guarantees that any point inside a δ neighborhood of \bar{q} will converge to \bar{q} . We will use Lyapunov's stability theorem to define a domain around \bar{q} where asymptotic stability is valid.

Theorem 5.2: Let \bar{q} be an equilibrium point for Eqn.4 and D be a domain containing \bar{q} . Let $V : D \rightarrow \mathbb{R}$ be a continuously differentiable function such that

$$V(\bar{q}) = 0 \quad (6)$$

$$V(q) > 0 \text{ for } q \in (D - \{\bar{q}\}) \quad (7)$$

$$\frac{dV(q)}{dt} < 0 \text{ for } q \in (D - \{\bar{q}\}) \quad (8)$$

then \bar{q} is asymptotically stable.

In this paper, we focus on systems of the form

$$\frac{dq}{dt} = A(q)(x_o - x(q)) = Ae \quad (9)$$

where x_o is a target pose and $A(q) : \mathbb{R}^r \rightarrow \mathbb{R}^n$ is a configuration-dependent linear map. The error e decreases linearly as the pose approaches the target.

Equilibrium points of Eqn.9 occur at

$$A(q)(x_o - x(q)) = Ae = 0 \quad (10)$$

If $\text{rank}(A) = r$, these occur at $x(q) = x_o$. We define the set of equilibrium points as the $(n-r)$ -dimensional self-motion manifold

$$\bar{\mathcal{Q}} = \{q \in \mathcal{Q} \mid x(q) = x_o\} \quad (11)$$

Consider an equilibrium point $\bar{q} \in \bar{\mathcal{Q}}$. We define a candidate Lyapunov function of the form

$$V(q) = \frac{1}{2}e'(Ae) \quad (12)$$

We know from Eqn.10 that $V(q) = 0$ at $q = \bar{q}$, thus satisfying the condition of Eqn.6.

However, if we define the domain D to be \mathcal{Q} , the requirements of Eqn.7 and Eqn.8 are *not* met: V is not uniquely 0 at \bar{q} . Rather, there is an entire set of equilibrium points $\bar{\mathcal{Q}} \subseteq \mathcal{Q}$ where $V = 0$. We address this issue by restricting V to a domain where \bar{q} is the exclusive minimum.

For illustration, we first consider the case where A has no functional dependence on q . We restrict the domain D of V to the r -dimensional hyperplane

$$D = \{q \in \mathcal{Q} \mid q = \bar{q} + Ae, e \in \mathbb{R}^r\} \quad (13)$$

Revisiting Eqn.7 and Eqn.8, we get $V(q) = 0$ and $\frac{dV(q)}{dt} = 0$ at $Ae = 0$ which occurs uniquely at $q = \bar{q}$ when V is restricted to D , thereby completing all the requirements for Lyapunov's theorem.

For the case where A does have a functional dependence on q , the Taylor series approximation of $A(q)$ in an n -dimensional ball $B_n(\bar{q})$ around \bar{q} gives

$$A(q) \approx A(\bar{q}) + \frac{dA}{dq}(q - \bar{q}) + \dots \quad (14)$$

We can then restrict the domain of V to the intersection of the r -dimensional hyperplane with $B_n(\bar{q})$

$$D = \{q \in \mathcal{Q} \mid q = \bar{q} + A(\bar{q})e, q \in B_n(\bar{q})\} \quad (15)$$

The intersection restricts the domain D to an r -dimensional ball $B_r(\bar{q})$ around \bar{q} (see Figure 4). Once restricted, the situation is identical to the previous case.

To summarize, we have shown that, for any point $\bar{q} \in \bar{\mathcal{Q}}$, there exists an r -dimensional ball $B_r(\bar{q})$ within which the conditions for Lyapunov's stability theorem hold.

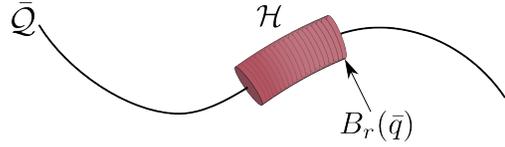


Fig. 4. A self-motion manifold in C-space. $r = 2$ and $n = 3$.

The evolution of Eqn.4 results in the projection $P(q_1) = \bar{q}$ for all $\bar{q} \in \bar{\mathcal{Q}}$, for all $q_1 \in B_r(\bar{q})$ if $q(0) = q_1$. Note that this is only valid for a self-motion manifold in isolation.

Let us now consider an open $(n-r)$ -dimensional ball $B_{\bar{\mathcal{Q}}}(\bar{q}) \subseteq \bar{\mathcal{Q}}$. Define $\mathcal{H}(B_{\bar{\mathcal{Q}}}(\bar{q}))$ as the manifold $B_r(\bar{q}) \times B_{\bar{\mathcal{Q}}}(\bar{q})$ embedded in \mathcal{Q} (see Figure 4).

\mathcal{H} has the following properties:

- 1) \mathcal{H} is n -dimensional.
- 2) \mathcal{H} contains $B_{\bar{\mathcal{Q}}}(\bar{q})$.
- 3) $P(q) \in B_{\bar{\mathcal{Q}}}(\bar{q})$ for all $q \in \mathcal{H}$ (conditional).

It is important to note that we have only described \mathcal{H} for a self-motion manifold in isolation. The third property of \mathcal{H} is only valid when $x(q)$ is closer to $x(\bar{\mathcal{Q}})$ than to any $x(k)$, for a self-motion manifold $k \subseteq (\mathcal{M} - \bar{\mathcal{Q}})$. I.e. the third property holds only when $\bar{\mathcal{Q}}$ is "chosen" by the projection operator.

C. Putting it all together

We will now bring the concepts developed thus far together to show coverage of \mathcal{M} by the sample-project method. Recall that, to show coverage of \mathcal{M} , we must show that the sample-project method places a sample inside any $B_m \subseteq \mathcal{M}$.

Section V-A showed that a configuration inside $x^{-1}(\mathcal{N})$ will project into \mathcal{S} , which consists of all self-motion manifolds that intersect B_m . Let us construct an n -dimensional manifold $\mathcal{UH}(B_m)$ by taking the union of the n -dimensional \mathcal{H} manifolds around every $\bar{\mathcal{Q}} \cap B_m$ for all $\bar{\mathcal{Q}} \subseteq \mathcal{S}$:

$$\mathcal{UH}(B_m) = \bigcup_{\bar{\mathcal{Q}} \subseteq \mathcal{S}} \mathcal{H}(\bar{\mathcal{Q}} \cap B_m) \quad (16)$$

\mathcal{UH} inherits the properties of \mathcal{H} , as well as the condition on the third property: A configuration $q \in \mathcal{UH}$ will project to a configuration inside B_m by the third property of \mathcal{H} if, for some $\bar{\mathcal{Q}} \subseteq \mathcal{S}$, $x(q)$ is closer to $x(\bar{\mathcal{Q}})$ than to any $x(k)$, for a self-motion manifold $k \subseteq (\mathcal{M} - \mathcal{S})$.

Lemma 5.3: Let $\mathcal{C} = (x^{-1}(\mathcal{N}) \cap \mathcal{UH})$. \mathcal{C} has the following properties:

- 1) \mathcal{C} is n -dimensional.
- 2) $P(q) \in B_m$ for all $q \in \mathcal{C}$.

Proof of 1st Property: $B_m \subseteq x^{-1}(\mathcal{N})$ by the second property of $x^{-1}(\mathcal{N})$. $B_m \subseteq \mathcal{UH}$ by the second property of \mathcal{H} . Both $x^{-1}(\mathcal{N})$ and \mathcal{UH} are n -dimensional. By Lemma 3.1, $x^{-1}(\mathcal{N}) \cap \mathcal{UH}$ is n -dimensional (see Figure 5). ■

Proof of 2nd Property: A $q \in (x^{-1}(\mathcal{N}) \cap \mathcal{UH})$ will project into \mathcal{S} because $q \in x^{-1}(\mathcal{N})$. Thus there exists a self-motion manifold $\bar{\mathcal{Q}} \subseteq \mathcal{S}$ such that $x(q)$ is closer to $x(\bar{\mathcal{Q}})$ than to any $x(k)$, for a self-motion manifold $k \subseteq (\mathcal{M} - \mathcal{S})$. This fact meets the condition required by the third property of \mathcal{H} .

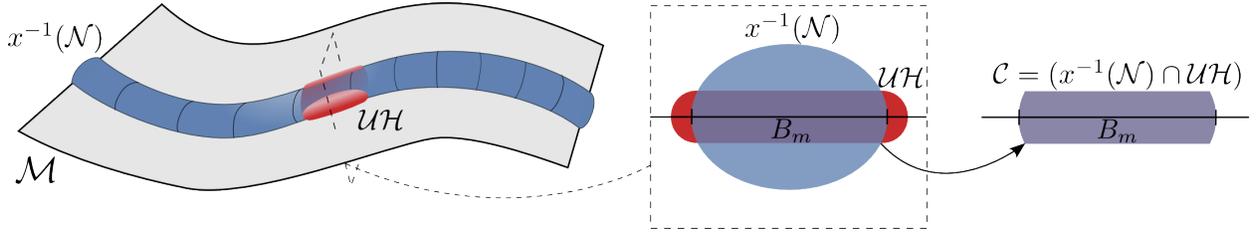


Fig. 5. An example of $\mathcal{C} = (x^{-1}(\mathcal{N}) \cap \mathcal{UH})$. $d = 1$, $r = 2$, and $n = 3$.

Because $q \in \mathcal{UH}$, and the third property of \mathcal{H} is valid, $P(q)$ must be in B_m . ■

Theorem 5.4: The sample-project method places a sample inside any B_m as the number of samples goes to infinity.

Proof: By the first property of \mathcal{C} , $\mu_n(\mathcal{C}) > 0$. Thus the probability of sampling a $q \in \mathcal{C}$ is greater than 0 and, as the number of samples goes to infinity, the probability of sampling a $q \in \mathcal{C}$ goes to 1. $P(q) \in B_m$ by the second property of \mathcal{C} . Thus we have shown that the sample-project method places a sample inside any B_m as the number of samples goes to infinity, which entails that the sample-project method covers \mathcal{M} . ■

VI. PROBABILISTIC COMPLETENESS OF RRT-BASED ALGORITHMS USING SAMPLE-PROJECT

We will use the fact that the sample-project method covers \mathcal{M} to prove that RRT-based methods that plan paths on \mathcal{M} are probabilistically complete. We focus on a class of RRT-based algorithms that plan with end-effector pose constraints, for example CBiRRT [1] and TCRRT [2]. These algorithms grow trees on \mathcal{M} by sampling near an existing node on \mathcal{M} and then projecting that sample to \mathcal{M} (see Figure 6). We will show that an RRT-based algorithm with the following properties is probabilistically complete.

- 1) Given a node of the existing tree, the probability of sampling in an n -dimensional ball centered at that node is greater than 0 and the sampling covers this ball.
- 2) The algorithm uses a projection operator with the properties of P to project samples to \mathcal{M} .

Define $\mathcal{M}_c \subseteq \mathcal{M}$ as the connected component of \mathcal{M} which contains the initial configuration. In order to be probabilistically complete, we must show that the algorithm will place a node in any m -dimensional ball in \mathcal{M}_c . Let the set of RRT nodes already generated by the algorithm (including the start and goal) be $N \subseteq \mathcal{M}_c$. A node of the tree is associated with a configuration q_n ; nodes will be referenced by their configuration. Let $B_n(q)$ for a configuration $q \in \mathcal{M}_c$ be an n -dimensional open ball centered at q .

Lemma 6.1: The algorithm covers $B_n(q_n) \cap \mathcal{M}_c$ as the number of samples goes to infinity.

Proof: Consider a $B'_m \subseteq (B_n(q_n) \cap \mathcal{M}_c)$. From Theorem 5.4, we know that any B'_m has an associated n -dimensional manifold \mathcal{C} such that all samples in \mathcal{C} project into B'_m . Both \mathcal{C} and $B_n(q_n)$ are n -dimensional and both contain B'_m . Thus, by Lemma 3.1, $\mathcal{C} \cap B_n(q_n)$ is n -dimensional. Because the algorithm samples in $B_n(q_n)$ with probability greater than

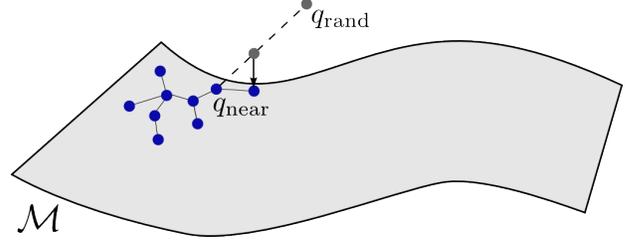


Fig. 6. An example of an RRT-based algorithm using the sample-project method. First a random sample q_{rand} is generated in the C-space and the closest node of the tree q_{near} is calculated. The algorithm then computes an intermediate configuration by stepping from q_{near} toward q_{rand} . If the distance between q_{rand} and q_{near} is smaller than the step size, q_{rand} is used as the intermediate configuration. The algorithm then projects this configuration to the manifold and adds the resulting node to the tree.

0 and the samples cover $B_n(q_n)$, it will generate a sample inside $\mathcal{C} \cap B_n(q_n)$ as the number of samples goes to infinity with probability 1. This sample will then project into B'_m by the second property of \mathcal{C} . Thus the algorithm projects a sample into any $B'_m \subseteq (B_n(q_n) \cap \mathcal{M}_c)$ as the number of samples goes to infinity. ■

We have shown that the algorithm covers \mathcal{M}_c locally, i.e. around existing nodes. We will now show that the nodes of the tree(s) cover \mathcal{M}_c as the number of samples goes to infinity. We will be using the series-of-balls argument in the subsequent proof. For a more detailed explanation of this argument see [12] and [13].

Theorem 6.2: The algorithm will place a node in any $B_m \subseteq \mathcal{M}_c$ as the number of samples goes to infinity.

Proof: For any $B_m \subseteq \mathcal{M}_c$, we can construct a series of open m -dimensional balls starting at some $q_n \in N$ such that subsequent balls overlap and the final ball overlaps with B_m . The overlapping regions between the balls are m -dimensional. By Lemma 6.1 a sample will be placed in any B'_m in the overlapping region if we sample in the B_n centered at the center of the previous ball. This sample will then be added to N as a node of the RRT. The algorithm is guaranteed to place a node within each overlap by induction, thus placing a node in any B_m as the number of samples goes to infinity. ■

Theorem 6.3: The algorithm is probabilistically complete.

Proof: By Theorem 6.2, as the number of samples goes to infinity any $B_m \subseteq \mathcal{M}_c$ will be sampled and a corresponding node will be added to N . As the radius of B_m goes to 0, N will approach \mathcal{M}_c in the limit. It follows that a path from start to goal will be found if one exists. ■

VII. DISCUSSION

We now discuss the implications of our proof for projection operators and mixed-dimensional constraint manifolds.

A. Projection Operators

Section V describes a projection operator that guarantees probabilistic completeness for RRT-based algorithms. The regularized Jacobian pseudo-inverse or Jacobian transpose iterative inverse kinematics methods can be used to perform the projection because they possess the requisite properties. Unfortunately, there are some common iterative methods which will *not* yield probabilistic completeness.

The null-space projection method [14] operates by using the null-space of the primary task (placing the end-effector in some pose) to satisfy secondary tasks such as collision-avoidance or balancing. For this method, $\frac{dq}{dt} = f(q(t))$ is:

$$\frac{dq}{dt} = \mathbf{J}^\# \dot{x} + (\mathbf{I} - \mathbf{J}^\# \mathbf{J}) \dot{q}_{\text{null}} \quad (17)$$

where $\mathbf{J}^\#$ is the generalized pseudo-inverse of the Jacobian, \dot{x} is the error in pose, and \dot{q}_{null} is the error in meeting a secondary objective. This type of projection attains optimal configurations by sliding along a self-motion manifold when $\dot{x} = 0$. The secondary task induces local minima on the self-motion manifold which attract configurations from the rest of the self-motion manifold. Thus, if a ball on the self-motion manifold does not contain a local minimum of the secondary task, configurations projecting to that ball may escape by sliding along the manifold. It follows this projection operator will not cover \mathcal{M} . Though using this projection operator in an RRT-based algorithm may yield an effective planner, it will not be probabilistically complete.

B. Mixed-dimensional Constraint Manifolds

The proof of coverage and probabilistic completeness assumed that the task constraint \mathcal{T} had a fixed dimensionality d . If we allow d to vary, then m (the dimensionality of \mathcal{M}) will vary as well. Since our proof of coverage used only local properties of \mathcal{M} , we can extend this proof to the case of varying m simply by applying the proof to each m -dimensional component of \mathcal{M} for every m . However, there is the case when a ball around a point on \mathcal{M} contains components of varying dimension. In this case, the ball can be split according to the dimensionality of its components and \mathcal{C} can be shown to exist for one of these components, thus guaranteeing a sample will be placed in this ball.

Though we can show that the sample-project method covers \mathcal{M} , the proof of probabilistic completeness for RRT-based algorithms only holds when \mathcal{M} is pure. The reason is that mixed-dimensional manifolds can be constructed such that all paths between two configurations must go through a narrow passage, which is of lower dimension than any component of the manifold. For instance, suppose \mathcal{M} were composed of two lines that intersect at some configuration q_p . To get from one line to the other, the algorithm would need to find a path which contained q_p . Yet there is 0 probability of generating q_p exactly, thus a path may never be found,

though one exists. This difficulty is not caused by the sample-project method and is not limited only to pose constraints. Rather this difficulty arises for all RRT-based planners when they must find a path through a lower-dimensional narrow passage. Coverage of \mathcal{M} does not entail probabilistic completeness in this case. In order to guarantee probabilistic completeness, the algorithm must be able to generate samples in these lower-dimensional narrow passages, as in [15].

VIII. CONCLUSION

We have presented a proof of probabilistic completeness for RRT-based algorithms when planning with end-effector pose constraints, regardless of their dimensionality. We first showed that the sample-project method used by these planners will cover the constraint manifold with samples. We then showed how this entailed probabilistic completeness. We also discussed the implications of the proof for projection operators and mixed-dimensional constraint manifolds.

IX. ACKNOWLEDGEMENTS

Dmitry Berenson was partially supported by Intel Labs Pittsburgh and by the National Science Foundation under Grant No. EEC-0540865. Thanks to Ross Knepper, Julius Ziegler, and David Handron for helpful discussions.

REFERENCES

- [1] D. Berenson, S. S. Srinivasa, D. Ferguson, and J. Kuffner, "Manipulation planning on constraint manifolds," in *ICRA*, 2009.
- [2] M. Stilman, "Task constrained motion planning in robot joint space," in *IROS*, 2007.
- [3] S. LaValle and J. Kuffner, "Rapidly-exploring random trees: Progress and prospects," in *WAFR*, 2000.
- [4] L. Sciavicco and B. Siciliano, *Modeling and Control of Robot Manipulators*, 2nd ed. Springer, 2000, pp. 96–100.
- [5] J. H. Yakey, S. M. LaValle, and L. E. Kavraki, "Randomized path planning for linkages with closed kinematic chains," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 6, pp. 951–958, 2001.
- [6] Y. Koga, K. Kondo, J. Kuffner, and J. Latombe, "Planning motions with intentions," in *SIGGRAPH*, 1994.
- [7] K. Yamane, J. Kuffner, and J. Hodgins, "Synthesizing animations of human manipulation tasks," in *SIGGRAPH*, 2004.
- [8] Z. Yao and K. Gupta, "Path planning with general end-effector constraints: using task space to guide configuration space search," in *IROS*, 2005.
- [9] J. Burdick, "On the inverse kinematics of redundant manipulators: characterization of the self-motion manifolds," in *ICRA*, 1989.
- [10] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006. [Online]. Available: <http://planning.cs.uiuc.edu/>
- [11] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2002.
- [12] P. Svestka, "On probabilistic completeness and expected complexity of probabilistic path planning," Department of Computer Science, Utrecht University, Utrecht, Netherlands, Technical Report UU-CS-96-20, May, 1996.
- [13] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *ICRA*, 2000.
- [14] S. Sentis and O. Khatib, "Synthesis of whole-body behaviors through hierarchical control of behavioral primitives," *International Journal of Humanoid Robotics*, vol. 2, pp. 505–518, December 2005.
- [15] K. Hauser and J. Latombe, "Multi-Modal Motion Planning in Non-Expansive Spaces," in *WAFR*, 2008.