IEEE Transactions on Robotics and Automation, Vol. 20, No. 2, p. 229–242, 2004

Measure Theoretic Analysis of Probabilistic Path Planning

Andrew M. Ladd Lydia E. Kavraki
Department of Computer Science
Rice University
Houston, TX, 77005
aladd|kavraki@cs.rice.edu

Abstract

*This paper presents a novel analysis of the probabilistic roadmap method (*PRM*) for path planning. We formulate the problem in terms of computing the transitive closure of a relation over a probability space and give a bound on the expected number iterations of PRM required to find a path in terms of the number of intermediate points and the probability of choosing a point from a certain set. Explicit geometric assumptions are not necessary to complete this analysis. As a result the analysis provides some unification of previous work. We provide an upper bound which could be refined using details specific to a given problem. This bound is of the same form as that proved in previous analyses but has simpler prerequisites and is proved on a more general class of problems. Using our framework we analyze some new path planning problems, $2k$-dof kinodynamic point robots, polygonal robots with contact, and deformable robots with force field control. These examples make explicit use of generality in our approach that did not exist in previous frameworks.*

I. Introduction

Planning a collision-free path for a rigid or articulated robot to move from an initial to a final configuration in a static environment is a central problem in robotics and has been the topic of extensive research over the last decade [1], [9], [11], [28]. The complexity of the problem is high and several versions of it have been shown PSPACE-hard [28]. Interesting applications and extensions of the problem exist in planning for robots that can modify their environments [12], [33] and flexible robots [3], [26], planning for graphics and simulation [24], planning for virtual prototyping [8], and planning for medical [39] and pharmaceutical [10] applications.

This paper concentrates on the analysis of PRM [23], [20], [36]. Since the early nineties, when PRM was invented, several researchers have reported on the advantages of the planner for various motion planning problems, particularily for robots with many degrees of freedom. Several variations of the method have been developed (e.g., [2]), several planners that bear resemblances with the original PRM have been introduced (e.g., [15]), and several extensions of the basic path planning problem have been solved with PRM-based methods (e.g., [33]). The experimental success of the planner has motivated many researchers to seek a theoretical basis for explaining its performance and relative successes in this direction have been reported, among others, in [22], [20], [36], [37], [21], [35], [14], [5], [29], [6], [15], [12], [7], [16]. This paper presents a further extension in this direction by using the mechanism of measure theory [4].

Algorithm 1 Operation of Basic PRM

```
Generate N points at random.
Connect the points with the local planner and obtain
  a directed graph G.
for each query of the form 'is there path from x to y?'
do
  Add x and y to point set and attach them to G.
  if a path from x to y lies in G then
    return the computed path.
  else
    return no path is found.
  end if
end for
```

The core operations of basic PRM are summarized in Algorithm 1. Points are generated at random from a probability distribution over the configuration space of the robot. The local planner is a computationally inexpensive heuristic for determining if there is a path between two given points. The local planner is usually not complete and will fail in some cases. However, if the local planner is sufficiently powerful and some path exists, a path can eventually be guessed. The relationship between the number of guesses, N, and the probability of finding a path between two given points is the subject of the analysis of PRM. In practice, PRM has proven to be suprisingly effective for certain applications and seems to be exploiting the fact that there are points that can be connected to many others.

Typically, there are two related questions that one asks about PRM operation.
1. How quickly can we find a path between some x and y?
2. How quickly can we find most of the paths?

The first question is usually addressed by isolating a particular path and showing an equivalent path can be constructed with a certain probability [21], [5], [36], [23], [37]. The second question is often dealt with by a space covering argument [22], [12], [15], [25].

The key difference in the treatment in our paper is the abstract reformulation of typical path planning problems to isolate the essential properties which allow us to analyze PRM. We capture the properties of a path planning task for a single path in terms of two simple constants. We relate these constants to previous analyses to suggest ways that they can be bounded in practice. Certain notions from an analysis such as [21] will be replaced with more general ideas. We will now describe some of the previous work that our treatment draws from.

Path Isolation Method

A common technique for analyzing `PRM` is to consider motion planning for point robots in open subspaces of $\mathcal{R}^d$. An example of this kind of analysis can be found in [21]. A single path is isolated and then analyzed by tiling with simple shapes or *buckets*. The analyses in [5], [21], [36] and [23] take this approach. The simple shapes used in these papers are ϵ-balls where $\epsilon > 0$ is related to the *path clearance*, the minimum distance of a path to the obstacles. The point choosing function is assumed to have distributions proportional to the volume of the balls. A result relating the probability of failure to the length of the isolated path and ϵ-clearance was given. An extension of this result was made for when ϵ varied along the path. Finally, this result shows that the probability of failure of finding a particular path goes exponentially to 0 as the number of milestones increases. The simplest of these results is related below.

Theorem I.1: [21] After generating N points in the free space $\mathcal{F} \subset [0,1]^d$ with the uniform distribution and joining them with straight lines, the probability of failure for finding a path of length L from point a to point b with path clearance ϵ was found to be

$$Pr[\text{FAILURE}] \leq \frac{2L}{\epsilon} e^{-\alpha_d \epsilon^d N},$$

where $\alpha_d = \frac{\omega_d}{|\mathcal{F}| \cdot 2^d}$ and ω_d is volume of the unit d-ball.

An extension to this technique can be made for small time locally controllable robots [38] such as car-like robots and tractor trailer robots. The property exploited is that for every point x and $\epsilon > 0$, there exists $\delta > 0$ such that any point within δ distance of x can be reached by taking a path that stays within the ϵ-ball around x. A path with ϵ clearance can thus be tiled with δ-balls. Again the probability of failure was shown to decrease exponentially as N increases. A proof for car-like robots that cannot drive backwards was also achieved in [37].

Space Covering Method

In [22], the notion of ϵ-goodness was described in an attempt to obtain the rate of total coverage of the free space. A space is ϵ-good if every point in the free space can 'see' more than an ϵ fraction of the space with the local planner. The spaces discussed were simply connected compact sets with measure 1. This paper showed how to bound the number of the guesses required to get an 'adequate' set of points in terms of ϵ. `PRM` has succeeded if an adequate set of points can be produced. It has been shown how to construct for any k an example of an 5/9-good polygonal space such that the space could not be completely covered by k points[40]. This disproved a conjecture in [22] that the number of points required to completely cover a polygonal configuration space was polynomial in the number of holes and the dimension. However, some evidence that the disconnected fraction of an ϵ-good space will decrease asymptotically to 0 was given.

In [12], an extension to ϵ-goodness using expansive spaces was considered. This work further formalized the notions of reachability and made use of measure in some sense. Concisely, for a connected set of points S, the $\beta - \text{LOOKOUT}(S)$ is the subset of S whose points 'see' using the local planner more than a β fraction of the set of points which can be 'seen' from S. A space is $(\alpha, \beta) - \text{EXPANSIVE}$ if the subset $\beta - LOOKOUT(S)$ is always larger than an α fraction of the measure of S for every connected subset S of the points reachable from any point in free space. Again this work provides a bound on the number of points in terms of α, β and ϵ required to generate a path as shown below.

Theorem I.2: [12] Let $\gamma \in (0,1]$ be a constant, M be a set of $2n+2$ points chosen independently and uniformly at random in free space $\mathcal{F}$ which breaks into connected components $\mathcal{F}_j$. Let G_j be the roadmap defined by M on $\mathcal{F}_j$. If

$$n \geq 8 \ln(8/\epsilon\alpha\gamma) + 3/\beta,$$

then with probability $1-\gamma$, G_j is connected.

This work was also useful when discussing probabilistic completeness of kinodynamic planners [15].

Key Concepts and Motivation

Our work began with the observation that the simple geometric shapes used to tile a path in a path isolation argument, e.g., ϵ-balls of [21], can be replaced with sets of strictly positive measure. These sets are not necessarily connected, open or even infinite. The probability distribution of point generation in Algorithm 1 can be replaced with a computable probability measure. The configuration spaces can be replaced with the more vague notion of state space without any explicit geometric assumptions. The predicate of reachability and the local planner are formalized with binary relations with the local planner being a subset of the reachability relation. The reachability is assumed to be transitive and both are assumed to have certain measurability properties.

Results and Outline

Our main result provides a bound on the expected number of points to generate to find a path from x to y using Algorithm 1 in terms of two space dependent constants n and p. The constant n is a kind of path length and the constant p is a generalization of path clearance. We also show that the probability of failure goes exponentially to 0. The result holds if and only if it there is positive probability of finding a random walk from x to y. In particular, if the probability of finding a random walk is zero then Algorithm 1 fails to find a path with probability one. We also provide a simple property to satisfy which is equivalent to probabilistic completeness, which in turn guarantees that some admissible n and p exist.

In the next section, we will formulate the problem intuitively. In Section III, we present preliminaries and notation. Following that, we prove the main result. The next three sections contain examples that show the power of our work in analysing path planning problems that could not be as easily addressed before. The final section is a discussion.

II. Problem Formulation

Our objective is to analyze the operation of Algorithm 1 for a given state space, set of valid paths, local planner and random point generation distribution. We seek to connect the value of N required to guarantee sufficiently low probability of error in PRM in terms of the local planner and the random point generation distribution.

Take a set X to be the free state space for a robot. This is taken to mean the entire set of distinct and allowable states the robot can assume. For example, this might be the configuration space with some extra information as to the state of the robot which is relevant to the planning problem, e.g., time, memory contents, velocity or the amount of fuel left. The path reachability relation is transitive, i.e., if x reaches y and y reaches z then x reaches z. This is a natural assumption which expresses what is intuitively understood by state space and path reachability. Note that symmetry and reflexivity are not enforced. If X encodes time, for example, path relations would be necessarily asymmetric.

The local planner can also be thought of as a binary relation over X. This relation, which we will call R, is not necessarily transitive. For example, a common local planner is a straight line local planner. It may be that a two segment path exists but a direct path passes through an obstacle.

A computational method that chains points together with a local planner has a hope of success if any valid path can be broken down into a finite sequence of states $x_1, ..., x_n$ such that $x_1 R x_2 \cdots x_{n-1} R x_n$. If $x_2, ..., x_{n-1}$ are present in the roadmap then a query of 'does x_1 reach x_n?' will be answered correctly. PRM can also return this sequence, from which the path can be reconstructed and executed [23].

PRM has a chance of finding a path in the case where the transitive closure of our local planner R, denoted $\bar{R}$, is the path reachability relation. Cases where the closure and the path relation do not agree will be discarded - the algorithm fails in these cases. For example, for a robot capable of moving in a planar box along any continuous curve, a poor choice of R would be a local planner such that (x, y) reaches (x', y') when $x < x'$ and $y < y'$. Some paths are unrealizable with such a local planner. We propose an abstract rephrasing of the general path planning problem with PRM. We are randomly determining membership in the transitive closure of the local planner.

III. Preliminaries

This section contains preliminaries, background, notation and describes the framework in which our analysis takes place.

A. *Measure Theory Review*

We first review abstract measure theory which is the language that our analysis is carried out in. We define measure space, product spaces, measureable functions, briefly discuss integrals and conclude by defining probability spaces. Our analysis of PRM proceeds by treating the state space as a probability space. These assumptions conviently incorporate various notions of sampling, topology, geometry and local planning.

A.1 Measure Spaces

A measure space (X, Σ_X, μ) consists of a set X, σ-algebra Σ_X over X and a measure μ. A σ-algebra over a set X is a collection of subsets of X that satisfies
1. $\emptyset \in \Sigma_X$,
2. for all $A \in \Sigma_X$, $X - A \in \Sigma_X$,
3. for any countable or finite indexing set I, and any collection of sets $A_i \in \Sigma_X$ indexed by $i \in I$, $\bigcup_{i \in I} A_i \in \Sigma_X$.

In other words, σ-algebras contain the empty set and are closed under complementation, finite and countable unions. The σ-algebra in a measure space is the set of the 'measurable' sets.

A measure μ is a function defined from Σ_X to $\mathcal{R}^{\geq 0}$ which assigns a 'size' to every measurable set. To be a well-defined measure, two properties must be satisfied:
1. $\mu(\emptyset) = 0$,
2. for any countable or finite indexing set I, and any collection of pairwise disjoint sets $A_i \in \Sigma_X$ indexed by $i \in I$,

$$\mu\left(\bigcup_{i \in I} A_i\right) = \sum_{i \in I} \mu(A_i).$$

A.2 Constructing σ-algebras

Let $\mathcal{S} \subseteq 2^X$ be any collection of subsets of X. By $\sigma(\mathcal{S})$, we denote the smallest σ-algebra containing $\mathcal{S}$. It can be shown that this construction is unique [4].

A.3 Borel Algebra

A topology for a space X is the collection of subsets that are 'open' in X. If X has a topology $\mathcal{T}_X$, then the Borel σ-algebra, sometimes written $\mathcal{B}(\mathcal{T}_X)$, is $\sigma(\mathcal{T}_X)$.

If $X = \mathcal{R}$ and Σ_X is the Borel σ-algebra for $\mathcal{R}$, a natural measure μ is defined by its operation on a closed interval $A = [a, b]$, $\mu(A) = |b - a|$. We sometimes refer to it as the usual measure for $\mathcal{R}$ and it can be extended by taking products to $\mathcal{R}^n$.

A.4 Product Spaces

Let (X, Σ_X, μ_X) and (Y, Σ_Y, μ_Y) be measure spaces. There is a natural product construction for a product measure space $(X \times Y, \Sigma_{X \times Y}, \mu_{X \times Y})$. The σ-algebra is chosen as $\Sigma_{X \times Y} = \sigma(\Sigma_X \times \Sigma_Y)$. The sets in $\Sigma_X \times \Sigma_Y$ are called 'rectangles'. It can be proved that every measurable set in the product can be written as a pairwise disjoint union of rectangles. This uniquely determines the action of $\mu_{X \times Y}$ by the following equation

$$\mu_{X \times Y}(C) = \mu_{X \times Y}\left(\bigcup_{i=1}^{\infty} A_i \times B_i\right) = \sum_{i=1}^{\infty} \mu_X(A_i) \cdot \mu_Y(B_i).$$

Given a measure space (X, Σ_X, μ_X), we can define the product space $(X^n, \Sigma_X^n, \mu_X^n)$ by the n-fold product construction as described above.

A.5 Measurable Functions and Integrals

Let (X, Σ_X, μ_X) and (Y, Σ_Y, μ_Y) be measure spaces. A function $f : X \to Y$ is called measurable if for every $B \in \Sigma_Y$, $f^{-1}(B) \in \Sigma_X$. If X and Y are topological spaces together with their Borel σ-algebras, then continuous functions between X and Y are measurable.

Let (X, Σ_X, μ_X) be a measure space. For any $A \in \Sigma_X$, the characteristic function $\chi_A : X \to \{0, 1\}$ is defined by the rule $\chi_A(x) = 1$ if and only if $x \in A$. Any measurable function $f : X \to \mathcal{R}$ is a simple function if it can be written $f(x) = c\chi_A(x)$ for some $A \in \Sigma_X$ and $c \in \mathcal{R}$. We can define integration over such a simple function by the rule

$$\int f(\cdot) d\mu_X = c\mu_X(A).$$

This rule holds for sums of n simple functions $f_1, ..., f_n$,

$$\int \sum_{i=1}^{n} f_i(\cdot) d\mu_X = \sum_{i=1}^{n} c_i \mu_X(A_i).$$

It is an important theorem of measure theory that any measurable function from a measure space X to $\mathcal{R}$ can be written as the limit of sums of simple functions. Furthermore this can be realized so that the limit and integral can be interchanged. This theorem is the underpinning for the integration theory that measure theory develops [34].

A.6 Probability Spaces

We say that (X, Σ_X, μ_X) is a probability space if μ_X is a probability measure. A probability measure μ satisfies the equation $\mu_X(X) = 1$.

B. *Basic Definitions of X, μ and R*

We now begin the analysis of Algorithm 1. To do this, we will first define the state space, X, the local relation R and probability measure μ which will represent the probability distribution of our point generating function. We will require that X forms a probability space with measure μ and that R be a 'measurable relation'.

As stated earlier, the set X will be the set of distinct and valid states the robot can assume. Let the set $\Sigma \subset 2^X$ be a σ-algebra for X. For example, a natural choice for this would be Borel algebra in the case where X has a topology [34]. Let function $\mu : \Sigma \to [0, 1]$ be a probability measure on (X, Σ). If α is the random variable indicating a point chosen from X at random by the sampler and A is a measurable subset of X, $Pr(\alpha \in A) = \mu(A)$. In short, (X, Σ, μ) is a probability space.

The local planner is described by a relation, R, over the set X. This relation will have the additional restriction that it is measurable, in other words $R \in \Sigma^2$. This is a natural assumption which, according to our understanding easily covers planning as we understand it today. The notation is given by the identity $xRy \Leftrightarrow (x, y) \in R \in \Sigma^2$. x reaches y is meant by xRy.

Another representation for R is as the characteristic function for the set R, which is more convenient for our purposes:

$$\chi_R(x, y) := \begin{cases} 0 & \text{for } (x, y) \notin R, \\ 1 & \text{for } (x, y) \in R. \end{cases}$$

The preimage of the above function is R, i.e., $R = \chi_R^{-1}(1)$.

C. *Transitive Closure*

Algorithm 1 seeks to sample X to learn facts about $\bar{R}$ inferred by using R. In this subsection we give a formal definition for $\bar{R}$ and prove that it is a measurable subset of X^2. This will be achieved by studying an iterated product construction. The objective is to show consistency in our definitions and to build some functions which will be used in the main result in Section IV.

The function χ_R is measurable and can be easily extended to n-ary analogues as follows. We will define a family of functions $f_n : X^n \to \{0, 1\}$ for which $f_n(x_1, ..., x_n)$ will be 1 if there is a path using R along the points $x_1, ..., x_n$ and 0 if there is no such path. In this notation, χ_R is f_2 and f_n is given by

$$f_n(x_1, ..., x_n) \mapsto \prod_{i=1}^{n-1} \chi_R(x_i, x_{i+1}).$$

In other words, $f_n(x_1, ..., x_n) = 1$ iff $x_1 R \cdots R x_n$.

Another useful version of this function will be written f_n^{xy} and defined as $f_n^{xy}(x_1, ..., x_n) := f_{n+2}(x, x_1, ..., x_n, y)$. $f_n^{xy}(x_1, ..., x_n)$ is 1 if and only if $x_1, ..., x_n$ is a path using R from x to y.

Finally, the transitive closure can be defined formally as

$$\bar{R} = \{(x, y) \in X^2 : \exists n, x_1, ..., x_n, f_n^{xy}(x, x_1, ..., x_n, y) = 1\}.$$

The following propositions are proved in the Appendix and are important to show that our definitions are consistent. The second proposition is used explicitly in the proof of the main result in Section IV.

Proposition III.1: $f_n(x_1, ..., x_n)$ is a measurable function.

Proposition III.2: f_n^{xy} is measurable for every $x, y \in X$.

Proposition III.3: $\bar{R}$ is a measurable subset of X^2.

D. *Guessing a Path at Random*

We will define a simple algorithm for path planning which tries to connect two points with a random walk. We will show that if this algorithm is probabilistically complete then so is Algorithm 1. Given a probability space (X, Σ_X, μ) and a local relation R, Algorithm 2 attempts to find a random walk from a point x to a point y using R and sampled from μ.

We will now define $\widehat{R}$ which will be a relation such that $x\widehat{R}y$ if and only if Algorithm 2 succeeds with probability

Algorithm 2 Random Walk Planner Given a Query (x, y)

Set $x_0 = x$.
Set $n = 0$.
loop
  Check if $x_n R y$, if so return $x_0, ..., x_n, y$ as the computed path.
  Generate x_{n+1} at random.
  Check if $x_n R x_{n+1}$, if not return **no path**.
end loop

greater than 0. We have defined three relations: R, the relation describing the local planner, $\widehat{R}$, the relation describing Algorithm 2, and $\bar{R}$, the transitive closure of R. It is obvious that the following chain of inclusions holds

$$R \subseteq \widehat{R} \subseteq \bar{R}.$$

The following lemma shows the consistency of our definition as well as providing an alternate construction for $\widehat{R}$.

Lemma III.4: $\widehat{R}$ is a measurable subset of X^2.

In the proof of this lemma, which can be found in the Appendix, the following observation is made.

Corollary III.5: $x\bar{R}y$ if and only if there exists n such that

$$\mu^n\left((f_n^{xy})^{-1}(1)\right) > 0.$$

Finally, we can prove a final result about completeness.

Theorem III.6: If $(x, y) \notin \widehat{R}$, then Algorithm 1 will find a path from x to y with probability 0. In particular, if $\widehat{R} \neq \bar{R}$ then Algorithm 1 will not be probabilistically complete for all queries.

For example, suppose $X \in [0, 1]^2$ and μ is the Borel measure on X. If due to constraints on R, finding a given path required guessing a point from a singleton set then for that query the probabability of success would be 0.

IV. Obtaining a Bound for N

The main result presented in this paper is a bound on the expected number of points needed to be generated in order to determine membership in $\bar{R}$. This bound is for any query $(x, y) \in \widehat{R}$. We know that if this assumption does not hold but $(x, y) \in \bar{R}$, then Algorithm 1 always fails on that query. The method of proof will be to reduce the problem of finding a path between two particular points to a standard problem in discrete probability; the following lemma will be used.

Lemma IV.1 (Coupon Collector[31]) To win a prize in a contest held by a breakfast cereal company, it is necessary to obtain at least one of each of n coupons in the boxes. The coupons are placed in the boxes according the uniform distribution, one per box. The expected number of boxes one must buy to get the prize is

$$E(N) = nH(n),$$

where $H(n)$ is the nth harmonic number. $H(n)$ is $\Theta(\log n)$.

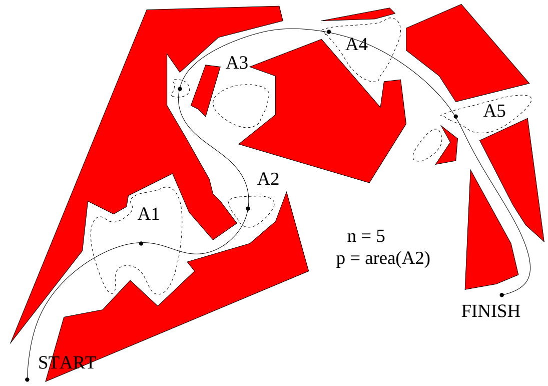

Fig. 1. A free path in state space and an illustration of the coupon buckets. Note that buckets A_3 and A_5 are disconnected.

The approach used in this paper is inspired by path tiling arguments such as [21] and [5]. In this line of argumentation, a covering of the path with balls is exhibited. This covering has the property that guessing a point inside each ball guarantees that the path has been found. This is very similar to a Coupon Collector game. The balls are the coupons in the game and can be thought of as buckets. If a randomly generated point lands in a bucket, then the associated coupon has been chosen.

To apply this proof scheme, the existence of certain buckets of strictly positive measure will be shown. These buckets will be such that guessing at least one point from each ensures that `PRM` has computed a path.

Theorem IV.2: If it is possible to find a path from x to y using Algorithm 2 with strictly positive probability, the expected number of random points needed to be generated for `PRM` to successfully compute a path from x to y using R, $E(N)$, satisfies the following inequality

$$E(N) \leq \frac{H(n)}{p},$$

for constants n and p.

Proof: Suppose $x\widehat{R}y$ for a given $x, y \in X$. We know by Corollary III.5 that $A = (f_n^{xy})^{-1}(1)$ is such that $\mu^n(A) > 0$. It follows that there is a rectangle $A_1 \times \cdots \times A_n \subseteq A$ such that $\mu(A_i) > 0$ for each i by the definition of μ^n. Notice that for any sequence $x_1, ..., x_n$ with $x_i \in A_i$, $xRx_1R\cdots Rx_nRy$ holds.

`PRM` will certainly succeed if a point from each A_i can be guessed. Since they each have positive measure this will eventually happen, however we would like to obtain a bound as well.

Each A_i can be thought of as a Coupon Collector bucket. For an illustration look at Figure 1. We will ignore points that land outside of the distinguished $A_1, ..., A_n$ in order to obtain an overestimation of $E(N)$. Also, we will assume that the A_i are disjoint - otherwise if a point is guessed which is in multiple buckets it can be randomly reassigned to a single bucket to obtain an overcount of $E(N)$.

Let $p = \min \mu(A_i)$ and conservatively take all the buckets to have measure p. Again this produces an overcount and we conclude that with probability at least np a point in at least one bucket is guessed.

We will define a sequence of random variables Y_i counting the number of guesses that are necessary, after the $(i-1)$th point has been guessed, in order to guess the ith point. The random variable T will count the number of guesses in any bucket required to obtain at least one guess in each bucket. It follows that $E(N) \leq E\left(\sum_{i=1}^{T} Y_i\right)$. Note that the Y_i are independent and identically distributed. Furthermore, the Y_i are independent of T. These observations allow us to conclude that $E\left(\sum_{i=1}^{T} Y_i\right) = E(T) \cdot E(Y_1)$.

By Coupon Collector we know that $E(T) = n \log n$ (since there are n buckets) and $E(Y_1) = 1/np$, the inverse of the probability of landing in some bucket. The final inequality follows. ∎

Corollary IV.3: After guessing N points, we can write the probability of not having guessed the path, P, as

$$P \leq n(1-p)^N.$$

Proof: This can be seen by applying the union bound, namely $Pr(\bigvee F_i) \leq \sum Pr(F_i)$. In our case, F_i says that bucket i is empty after N point guesses. $Pr(F_i) \leq (1-p)^N$ and there are n such F_i's. ∎

The limiting part of bound of Theorem IV.2 is p. Our analysis allows for the use of sets which may not have simple geometric shapes thus extending the palette of sets of positive measure that we could use to find a lower bound for p. An alternate way of viewing p is as a generalization of path clearance.

Synopsis and Examples

Up to this point in the paper we have developed and proved Theorem IV.2. We continue by giving several examples that can be analyzed in a novel way by applying the main result and the techniques used to derive it.

1. In the first example, we consider $2k$-dof kinodynamic robots. These are point robots in $\mathcal{R}^k$ with the k velocities as part of the state space. We give a short proof which illustrates a general technique for proving $\widehat{R} = \bar{R}$. Also we show that by using disjoint, oddly shaped buckets, we can improve our bounds over methods which use only balls.
2. In the second example, we describe a problem consisting of a polygonal robot moving in a polygonal workspace. Contact at the boundary is permitted and consequently computing path reachability with a probabilistic planner is a challenging task. To construct an appropriate local planner and sampling distribution, we proceed by introducing a general construction of measures for an important class of topological spaces. The general construction is then applied to obtain an admissible planner.
3. The third example gives a construction for path planning with a deformable robot with no fixed parametrization under energy constraints. We propose a solution using a partial order of subdivision spaces and argue that for any path, a similar path appears at some subdivision level. We then show how a planner can be constructed that plans over all subdivision levels. Although the example is not worked out to the last detail and questions of efficiency are not considered, we show how we can build path planners for non-manifold spaces with no fixed parametrization, a result which is first suggested in this paper.

V. $2k$-dof Kinodynamic Robots

We begin by summarizing the known single path analysis of a k-dof holonomic robot and then discuss issues surrounding the extension to $2k$-dof kinodynamic robots.

The workspace for this example will be a k-manifold $W \subset [0,1]^k$. The state space is the workspace, $X = W$. The random sample function has the distribution induced by the Borel measure on $[0,1]^k$ normalized to be a probability measure.

Consider a fully holonomic robot operating in this workspace with a local planner that connects points with a straight line. This analysis closely parallels [21] and [5].

The buckets (as in Theorem IV.2) for k-dof point robots can be constructed and used to get more explicit bounds on $E(N)$. Let $\mu(B_\delta(\cdot))$ be the measure of an open δ-ball in X.

Proposition V.1: Suppose we have some path with ϵ *path clearance* and length L. Then

$$E(N) \leq \frac{H(2L/\epsilon)\mu(X)}{\mu(B_{\epsilon/2}(\cdot))}.$$

Proof: Let $\gamma : [0,1] \to X$ be a simple path from x to y. Let $\epsilon > 0$ be the path clearance and suppose γ has length L. We see in [5] that γ can be tiled with balls of radius $\epsilon/2$ and that $2L/\epsilon$ balls are sufficient. We can set $n = 2L/\epsilon$ and $p = \mu(B_{\epsilon/2}(\cdot))/\mu(X)$ and then the result follows from Theorem IV.2. ∎

Estimating $E(N)$ is made more difficult by considering velocities, making the problem a much more complex $2k$-dof problem. Some progress in this direction has been made [15], [25]. We show that `PRM` succeeds but that ball tiling arguments as in [21] are inappropriate for this problem. The shapes of the buckets we need are disconnected and dependent on the input points.

The workspace once velocities are considered is a k-manifold $W \subset [0,1]^k$ and the state space of the robot is $X = W \times (-1,1)^k$, which encodes position and velocity. The robot can be controlled by applying a constant acceleration in the range $(-1,1)$ for a constant non-zero time period. In 1-D, the local planner that we use tries to connect position (x_1, v_1) with (x_2, v_2) by taking

$$\Delta t = \tfrac{2(x_2 - x_1)}{v_1 + v_2}, \quad a = \tfrac{v_2 - v_1}{\Delta t}\ .$$

In k-D, we solve each dimension independently and cases where either acceleration is too large, acceleration is zero, singularities arise or the time increment is non-positive are not considered to be solutions. We refer to this local planner as R in the following. We note that given a C^1 path between two points, a path made of piecewise constant second derivatives which arbitrarily well approximates the first path can be found [34].

The claim we will establish is that for a given primitive path, xRy, we can move the endpoints around in open

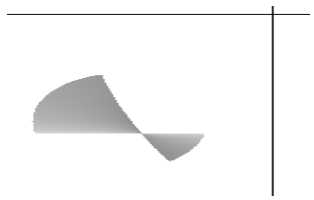

Fig. 2. The grey area denotes the set of points b for which $aRbRc$ for given a and c for 1-d robot. The x-axis is position and the y-axis is velocity.

neighbourhoods without breaking the path. Together with closure of open sets under intersection, we show this fact is enough to conclude that $\widehat{R} = \bar{R}$.

Let $Y = \{(x, v, y, w) : (x, v)R(y, w)\}$. For every $\bar{x} = (x, v, y, w) \in Y$, $|x-y| > 0$ and $|v \pm w| > 0$, so $\exists \epsilon > 0$ such that $B_\epsilon(\bar{x}) \subset Y$. It follows that Y is 4-manifold.

The point guessing distribution, μ, that we use has positive measure on open sets. Suppose $(x_1, v_1)R(x_2, v_2)$ $R(x_3, v_3)$. There exists $\epsilon > 0$ such that $(x_2', v_2') \in B_\epsilon(x_2, v_2)$ is such that $(x_1, v_1)R(x_2', v_2')R(x_3, v_3)$. Since $\mu(B_\epsilon((x_2, v_2))) > 0$, we have a probabilistically complete path planner, i.e., $\bar{R} = \widehat{R}$, and Theorem IV.2 applies.

In this example, n and p also depend on the input points. In Figure 2, we can see a possible solution space for a planning problem with a single intermediate point in 1-D, the vertical axis being velocity and horizontal axis being position. It was computed by sampling the kinematics equation. The shape depends heavily on the start and finish points and is disconnected. The measure of the set, however, is a significant fraction of the measure of the smallest disc which encloses all of the points. An analysis based on ball shaped buckets would not be able to take this properly into account. By explicitly taking into account the shape of the buckets we can obtain tighter bounds.

VI. Planning with Contact

In this section, we describe a class of path planning problems for simple polygonal robots operating in a polygonal workspace. The boundary of the robot is allowed to make contact with the boundary of the obstacles. Certain instances of path planning tasks in this class can be made to require that the only paths between certain pairs of configurations necessarily touch the boundary. It can be shown that the usual measure on the set of transformations fails to solve this problem. One has only to consider an example which requires passing through some point in a set of boundary points. Since the measure of the boundary is zero, no solution can be found by the usual sampling techniques. This problem can be fixed by taking a different measure. We give a general construction for a measure that allows us to show that path reachability for any finite CW complex can be computed with `PRM`. The random sampling techniques are similar in spirit to those discussed in [17], [19], [18]. The resulting planner will be proved to be complete which to our knowledge is a novel result. Our analytical technique is general and supports further extensions into 3-d, perhaps by using the methods from [17], [19], [18] or a similar approach. The reasons why we chose CW complexes are explained further. We end this section by showing how the polygonal robot problem we suggested can be related to this framework and sketch a full solution.

A. Polygonal Robot with Contact

Let P be a simple polygon with vertex set $\{p_1, ..., p_n\}$. Let $T = S^1 \times \mathcal{R}^2$ be the set of transformations that embed P into $\mathcal{R}^2$. We interpret $(\theta, a, b) \in T$ to be a rotation about the origin by θ followed by translation by (a, b). $r \in T$ applied to P transforms P into $r(P)$.

Let $\sigma_1, ..., \sigma_k$ be a set of triangles in $[0,1]^2$ such that they have pairwise non-intersecting interiors. Let $O = \bigcup \sigma_i$ be the polygonal obstacle space and let $Int(O)$ be the interior of those obstacles. We can now define the free configuration space, X, by

$$X = \{r \in T : r(P) \cap Int(O) = \emptyset \text{ and } r(P) \subset [0,1]^2\}.$$

X will take the induced topology from T. The path reachability relation is defined in the usual way, by existence of a continuous map $\gamma : [0,1] \to X$.

The task is to find a probability measure μ and appropriate local planner R so that `PRM` computes the desired path reachability relation. Observe that under the usual measure, the set of all contact configurations has measure 0. We also require that R be easily computed.

B. Finite CW Complexes

Many `PRM` analyses in the literature have been for strictly manifold configuration spaces where the reachability relation was path reachability. We will describe how `PRM` can be used to calculate path reachability for a certain class of non-manifold configuration spaces with a useful structure called *finite CW complexes*. The name CW is due to J.H.C. Whitehead and refers to **C**losure-finiteness and **W**eak topology, two defining properties of CW complexes [32].

The importance of CW complexes arises in part from being a very general class of topological spaces for which there is a systematic way of computing homologies. In particular, a CW complex can be decomposed into a set of cells, where each cell is a connected k-manifold for some k. We give a way to extend the computation of path reachability in each individual cell to path reachability over a space composed of many cells.

The simplicity of CW complexes derives from their definition. They are defined inductively. We call X_0 the 0-*skeleton*. X_0 is a finite union of closed 0-balls (points). Given a k-skeleton X_k, we construct the $k+1$-skeleton by considering a finite union of closed $k+1$-balls, Z_{k+1}. Each ball $B \in Z_{k+1}$ is called a cell of dimension $k+1$. We observe that Z_{k+1} is a $k+1$-manifold with boundary and denote its boundary by $Bd(Z_{k+1})$. Let $f : Bd(Z_{k+1}) \to X_k$ be continuous. X_{k+1} is the space obtained by identifying points of $X_k \cup Z_{k+1}$ according to the map f. This identification can also be thought as defining a quotient map

$\nu : X_k \cup Z_{k+1} \to X_{k+1}$. For every k, a space X_k constructed in this fashion is a finite CW complex. A finite CW complex is clearly compact.

For example, to obtain the CW complex for a torus, we begin with a point a and two lines l and m to this point to obtain two loops. Then a disc is attached to the loops by winding it along l, then along m, then along l in the opposite direction and then along m in the opposite direction.

In Section IV, we gave Theorem III.6, a necessary and sufficient condition for completeness. We will now give a sufficient condition for completeness which has natural and convenient statement for CW complexes and further illustrate its use by applying to the contact robot problem we described. We choose to give only a sufficient condition in this case as it makes the treatment much simpler.

Theorem VI.1: Let X be a finite CW complex with decomposition $X_0 \subset \cdots \subset X_n$. Suppose we are given a measure μ and a local planner R such that for every cell B,
1. if $A \subset B$ is open and non-empty in B, then $\mu(A) > 0$,
2. if xRy with $y \in Int(B)$, there there exists $A \subset Int(B)$ such that A is non-empty and open in B and $y' \in A$ implies xRy',
3. R is symmetric and $\bar{R}$ contains path reachability for B.
It follows that a `PRM` planner using μ and R is probabilistically complete for computing path reachability in X.

A proof of this theorem is given in the Appendix.

It is interesting to note that path reachability is not computable in general using the scheme we described without the caveat that a CW complex is finite. A counterexample where the transitive closure of the local relation is not path reachability can be formulated. It is a spiral of boxes which tiles the unit rectangle. The center point of the spiral is singular and is disconnected in the transitive closure of the local planner we gave. The unit rectangle has a much simpler finite CW complex which avoids this problem. Oddly enough, many infinite CW complexes do not have this problem, for example the Hawaiian Earring [32] or a non-compact k-manifold. A finer distinction can be made by requiring the complexes to be such that any two path-connected points have a connecting path which passes through finitely many cells.

C. A Solution to the Contact Robot Problem

Now that we have developed Theorem VI.1, we return to the problem described at the beginning of this section. To provide μ, an appropriate probability measure, and R, a computable local planner, we will show that X, the configuration for a polygon robot in contact, is a finite CW complex and then apply Theorem VI.1. To do this, we will provide explicit parametrizations for each of the subspaces and define a straight line planner and measures. The main result of this section is stated as follows.

Theorem VI.2: X is a finite CW complex. There is an explicit measure and local planner for X that satisfy the conditions of Theorem VI.1.

Once we have proven this theorem, we can invoke Theorem VI.1 and Theorem IV.2 to prove that the constructed planner is probabilistically complete. Since the proof requires an explicit construction of parametrizations of a decomposition of the space, the construction engenders an implementation of this planner. Although we give a more complete version of the full argument in the Appendix, we will sketch the main steps of the proof. To proceed, we will first define contact constraints.

The boundary of the obstacle space, O, can be decomposed in a set of vertices $v_1, ..., v_l$ and set of edges $e_1, ..., e_m$. Similarily, we have vertices from the robot $p_1, ..., p_n$ and edges $s_1, ..., s_n$. Let $r \in X$ be a configuration of the robot. The contact topology of r is a list of constraint pairs of the form (v_i, s_j) or (e_i, p_j). We will define $\mathcal{P}$ as the collection of all subsets of constraint pairs. For any $x \in X$, we write $\alpha(x)$ as a unary relation to say that x satisfies constraint α. We also say that $\alpha \preceq \beta$ if $\alpha \subset \beta$ and $\alpha \prec \beta$ if $\alpha \subset \beta$ and $\alpha \neq \beta$.

For a given α, we will outline explicit parametrizations for spaces of candidate points satisfying α. This is written in the Appendix. We can now describe the sampling function we will use.

Algorithm 3 Sampling for Contact Robot (μ)

```
loop
    generate a random constraint α such that |α| ≤ 3.
    from the parametrization for points satisfying α, generate a candidate point, x, using uniform sampling.
    if x ∈ X then return x.
end loop
```

The local planner is defined similarily.

Algorithm 4 Local Planner for Contact Robot: xRy

```
find α maximal such that α(x) and α(y).
find the straight line l in the parametrization for α between x and y.
check if l ⊂ X, if so, return true.
otherwise, return false.
```

The subspace Z_α consists of all points which properly satisfy constraint α. A formal definition is

$$Z_\alpha = \{x \in X : \alpha(x) \text{ and } \beta(x) \Rightarrow \beta \preceq \alpha\}.$$

Now, we will prove some technical lemmas. The first lemma states that every point of X appears in exactly one Z_α. The second lemma requires a more intricate argument and must be proved case by case. It states that the non-empty Z_α are k_α-manifolds.

Lemma VI.3: The Z_α are pairwise disjoint and their union is X.

The proof of this lemma follows directly from the definition of Z_α and is related in the Appendix.

Lemma VI.4: For every α, either $Z_\alpha = \emptyset$ or there is an integer $0 \leq k_\alpha \leq 3$ such that Z_α is a k_α-manifold.

The proof of this lemma is dealt with in the Appendix.

Theorem VI.2 requires first proving that X is a finite CW complex. This is done by using cell structure defined by the Z_α. Lemma VI.4 states that Z_α are manifolds and Lemma VI.3 states we have exhausted the space. The cell structure is defined for $0 \leq d \leq 3$,

$$X_d = \bigcup_{\alpha: k_\alpha \leq d} Z_\alpha.$$

Although the argument we relate for showing the probabilistic completeness of the planner we constructed for solving the polygon robot with contact problem is intricate, it is greatly simplified by being able to appeal to Theorems IV.2 and VI.1. The resulting planner, as described by Algorithms 1, 3 and 4, is also algorithmically concise. Although it is intuitive that this algorithm would be probabilistically complete, proving that it is the case is well beyond the scope of the earlier completeness proofs for PRM that appeared in the literature.

VII. Deformable Robots

In this section, we consider motion planning with deformable robots controlled by force fields. This section will sketch how to show probabilistic completeness of the path planner. There will be little emphasis on the control and simulation of parametric deformables, an interesting topic on its own.

An example of planning for deformable robots developed in [27] is that of a deformable sheet manipulated in 3-D workspace under energy constraints. The manipulation can be abstractly modeled as an energy field. An illustration is given in Figure 3.

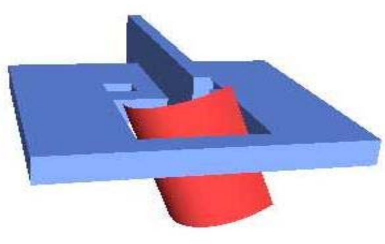

Fig. 3. An example of a deformable robot (curved surface) that has to pass through a hole in a polyhedral obstacle.

The robot we consider in this section will be a deformable curve operating in a k-dimensional compact workspace. It will be controlled by an external force field.

For the sake of simplicity, suppose the configuration space is the set of all C^2 curves embedded into a k-manifold workspace $W \subset [0,1]^k$ which satisfy some constraints on total deformation energy and local strain energy. The state space X consists of embeddings of the curve into the workspace, W, together with a C^1 velocity field on the curve. The robot can be controlled by applying C^0 force fields to the curve. Previous work on this problem took a fixed parametrization and this is an approximation which might not be able to express some of the paths in X. In this section, we will show that Algorithm 1 can be applied over all parametrizations.

We can subdivide the curve recursively (say in two pieces). This will form a partial order $\mathcal{L}$ of subdivision topologies. For each $\lambda \in \mathcal{L}$, we have a state space X_λ which is an m-manifold for some m which represents the curve's constrained deformation, embedding and velocity field in terms of a finite parameter set (where the curve is obtained by interpolation). To each $\lambda \in \mathcal{L}$, we assign a probability $p_\lambda > 0$ such that $\sum_{\lambda \in \mathcal{L}} p_\lambda = 1$. Also, an operator $\vee$ on every pair $\lambda, \lambda' \in \mathcal{L}$ can be defined so that $\lambda \vee \lambda'$ is the simplest common subdivision topology.

For two states $x, y \in X$, suppose there is a path between them. We will now sketch an approximation scheme for the path, discuss what kind of properties the local planner must have and show how we can compute the path with PRM. More specifically, we construct X' with an associated measure and local planner R such that PRM succeeds and implies paths in X. We note that X is not finitely parametric and that X' is not a manifold.

We rely on several reasonable assumptions. The subdivision scheme we are using must be of the type where the curve represented by some subdivision topology and parameters must be the limit of the subdivision process. The family of curves and primitive paths must also be sufficiently rich to approximate any given curve arbitrarily well when taken under finite composition, i.e., $\bar{R}$ is path reachability. Finally, we assume that queries are made with representable pairs (x, y).

Let R_λ be the local planner which connects points in X_λ. We must first show that X_λ with its probability measure μ_λ and with local planner R_λ satisfy the conditions for Theorem IV.2. Recall that X_λ is an m-manifold. For any $z_1, z_2, z_3 \in X_\lambda$ such that $z_1 R_\lambda z_2 R_\lambda z_3$, we define Y as the set of points $z_2' \in Y$ where $z_1 R_\lambda z_2' R_\lambda z_3$. It is now sufficient to show that Y is also an m-manifold. If we can conclude that $\mu_\lambda(Y) > 0$, then it follows that $p > 0$ (in the sense of Theorem IV.2) for any path with a finite number of points. We will assume for the moment that this assertion holds.

We construct the state space $X' = \bigcup_{\lambda \in \mathcal{L}} X_\lambda$ with probability measure taken by the product σ-algebra and measure constructions. Measures are weighted for each λ by p_λ. The local planner for points z, z' (with subdivision λ and λ' respectively) works by reinterpreting z and z' as points in $X_{\lambda \vee \lambda'}$ and using its corresponding local planner. This constructs a probability space by the product construction. Furthermore, it follows that PRM is probabilistically complete on X' with local planner R.

Suppose $\gamma : [0,1] \to X$ is the path between x and y and this path has ϵ clearance. Since the subdivisions can generate arbitrarily good approximations to points in X, there exist points $x_1, ..., x_n \in X'$ such that $x R x_1 R \cdots R x_n R y$ and the new path is within ϵ of γ.

This observation holds under the assumption that the set of valid center points for a length two path in R with fixed endpoints has positive measure. This is a reasonable fact that might be proved using the open set technique from the section on $2k$-dof robots. This assumption will hold under a variety of reasonable conditions. We do not give a proof to avoid introducing unwieldly formalism.

We have shown that, under reasonable assumptions, given a path for our deformable robot we can construct a path which is within $\epsilon > 0$ that can be found using PRM without fixing a parametrization a priori. This shows that a generic path planner could be constructed for this problem and that the probability of failure the planner would go to 0 exponentially in the number of guesses. Since the approximation space we constructed is not a manifold, we also note that we succeeded in showing path planning results in non-manifold spaces without sacrificing the aspects of PRM that make it desirable to implement in practice. Finding a path in the approximation space implies a path in X. This kind of analysis was not possible with previous frameworks.

VIII. Discussion and Future Work

In this paper, we reformulated the robot path planning problem in terms of probability spaces, measures and computation of the transitive closure of a given local relation. We showed that if it was possible to guess a path between two given points at random, then n sets of strictly positive measure existed so that guessing at least one point in each set would produce some path between these points. This allowed us to bound the probability of failure in terms of n and p, where p is the measure of the smallest bucket. We also used Lemma IV.1 to find a bound for the expected number of points we need to guess in order to find some path between the query points, if that path exists. Since the assumptions we used for our proofs were quite weak, our treatment has bearing on using PRM for planning problems in practice.

It is interesting to note that when there are buckets with variable sizes, Coupon Collecting is limited by the smallest bucket. This bears a strong similarity to results on narrow corridors published in [13]. While this is not a proof that PRM is limited by small p, it is consistent with previous claims to this effect. Certainly one can easily construct examples where a point must be guessed from a very small bucket and then argue a lower bound inversely proportional to p, the size of this small bucket.

In our first example, $2k$-dof kinodynamic robots, we echoed the results on kinodynamic planning of [15] and gave a rigorous proof using a path isolation analysis. In this example, we used buckets which were possibly disconnected and had shapes which were not at all similar to that of a ball.

The second example presented in Section VI-A motivates our framework by showing the need to reason explicitly about the probability measure. This was used to solve the problem of computing path reachability on finite CW complexes. We were motivated by polygon manipulation with contact. The defined spaces are *not* manifolds and we used an unusual probability measure to guarantee a solution. This work has direct applications to planning with contact and for certain kinds of manipulation. Some interesting and direct extensions of this work can be made by adding dynamics and planning for collisions or for manipulation by 'pushing'.

Our treatment of the second example showed constructions which are useful for reasoning about planning with contact and for manipulation. The combinatorial structure of polyhedra in contact is quite complex and constructing the CW complexes for these contacts explicitly can be quite difficult particularily when multiple contacts are allowed. Realizing such construction in practice can still be quite difficult.

In the final example presented in the paper, we showed a probabilistically complete path planner for a deformable curve robot which did not make use of a fixed parametrization. This is a powerful extension of [3] and further motivates using probability spaces for the analysis of PRM. We constructed an approximation space which was easier to represent and implied paths in the original space. The construction that we proposed was complete in the sense that any path occurred in some approximation space. This technique may have additional applications for certain problems where incomplete approximation spaces could be used.

A. Space Coverings

Our approach used a path isolation argument. A space covering argument for PRM analysis is another approach that has been used. It would be interesting to try to relate our analysis to [12] and recast it in terms of space coverings. We know that for any given pair of path connected points, the probability of not finding a witness to their connectivity decreases as the number points we guess with PRM grows. It seems natural that the measure of the missing part of $\mu^2(\bar{R})$ can be made arbitrarily small by guessing more points. More formally we have the following.

Conjecture VIII.1: Let (X, μ) be a probability space and let R be a local relation such that the conditions for Theorem IV.2 are satisfied. We define R_N as the relation computed by PRM after N guesses. Prove that for every $\epsilon > 0$, there exists N such that

$$E[\mu^2(\bar{R}) - \mu^2(R_N)] < \epsilon.$$

B. Polygon Manipulation

Combining some of the ideas we discussed in the examples can lead to many challenging problems in path planning. We end this paper with such a problem and some conjectures about it. This problem is very interesting for practical applications in assembly and part manipulation.

Conjecture VIII.2: Let $P_1, ..., P_n$ be simple polygons in the plane and let the workspace be $W = [0,1]^2$. Suppose these polygons have masses $m_1, ..., m_n$ distributed uniformly. The polygons can contact each other and will collide using an impulse model (such as [30]). A polygon cannot leave the workspace. Let F be a class of force fields. A force field f can be applied to the workspace and the polygons will move accordingly. They may or may not converge to a fixed configuration.

For a given choice of F:

1. For some $n > n_0$, there is a choice of polygons and target configuration y such that for any $f \in F$, there exists

an initial configuration x such that x with f applied to it does not converge to y.
2. For some $n > n_0$, there is a choice of polygons, starting configuration x and target configuration y such that for any $f \in F$, x with f applied to it does not converge to y.
3. For any n, choice of polygons, start configuration x and end configuration y, the existence of a contact-free holonomic path from x to y implies that there is a sequence of force fields $f_1, ..., f_{k+1} \in F$ and times $t_1, ..., t_k$ such that applying f_i for time t_i and then ending with f_{k+1} applied for ever will bring configuration x to configuration y. Furthermore, a `PRM` planner can be constructed for this task.

The first two parts argue that a result such as [27] cannot be obtained for this system and the third claims that `PRM` could be used for this problem.

Acknowledgments

Andrew Ladd is partially supported by FCAR, NSF IRI-970228 and a grant from the Whitaker Foundation. Work on this paper by Lydia Kavraki was supported by NSF Award IRI-970228, a Whitaker grant, an ATP Award and a Sloan Fellowship. Andrew Ladd would like to thank Gregory Dudek for introducing him to this topic, to robotics in general and for many discussions on randomized path planning.

References

[1] P. K. Agarwal, L. E. Kavraki, and M. Mason. *Robotics: The Algorithmic Perspective.* A K Peters, Natick, MA, 1998.

[2] N. Amato, B. Bayazit, L. Dale, C. Jones, and D. Vallejo. Obprm: An obstacle-based prm for 3d workspaces. In P. Agarwal, L. Kavraki, and M. Mason, editors, *Robotics: The Algorithmic Perspective*, pages 156–168. AK Peters, 1998.

[3] E. Anshelevich, S. Owens, F. Lamiraux, and L. Kavraki. Deformable volumes in path planning applications. In *IEEE Int. Conf. Robot. & Autom.*, pages 2290–2295, 2000.

[4] P. Billingsley. *Probability and Measure.* Wiley-Interscience, 1995.

[5] R. Bohlin. *Robot Path Planning.* PhD. Thesis, Chalmers Univ. of Technology, 2002.

[6] R. Bohlin and L. Kavraki. Path planning using lazy prm. In *IEEE Int. Conf. Robot. & Autom.*, pages 521–528, 2000.

[7] R. Bohlin and L. Kavraki. *A Randomized Algorithm for Robot Path Planning Based on Lazy Evaluation*, pages 221–249. Kluwer Academic Publishers, 2001.

[8] H. Chang and T. Y. Li. Assembly maintainability study with motion planning. In *Proc. IEEE Int. Conf. on Rob. and Autom.*, pages 1012–1019, 1995.

[9] B. Donald, K. Lynch, and D. Rus. *New Directions in Algorithmic and Computational Robotics.* A K Peters, Boston, MA, 2000.

[10] P. W. Finn and L. E. Kavraki. Computational approaches to drug design. *Algorithmica*, 25:347–371, 1999.

[11] K. Gupta and A. P. del Pobil. *Practical Motion Planning in Robotics.* John Wiley, West Sussex, UK, 1998.

[12] D. Hsu. *Randomized Single-Query Motion Planning In Expansive Spaces.* PhD thesis, Department of Computer Science, Stanford University, 2000.

[13] D. Hsu, L. Kavraki, J. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In P. Agarwal, L. Kavraki, and M. Mason, editors, *Robotics: The Algorithmic Perspective*, pages 141–154. A K Peters, 1998.

[14] D. Hsu, L. E. Kavraki, J.-C. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In e. a. P. Agarwal, editor, *Robotics: The Algorithmic Perspective*, pages 141–154. A.K. Peters, Wellesley, MA, 1998.

[15] D. Hsu, R. Kindel, J. Latombe, and S. Rock. Randomized kinodynamic motion planning with moving obstacles. In *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2000.

[16] D. Hsu, R. Kindel, J. Latombe, and S. Rock. Randomized kinodynamic motion planning with moving obstacles. *Int. J. Robot. Res.*, 2001.

[17] X. Ji and J. Xiao. On random sampling in contact configuration space. In *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2000.

[18] X. Ji and J. Xiao. On automatic generation of high-level contact state space. *Int. J. Robot. Res.*, 20(7):584–606, 2001.

[19] X. Ji and J. Xiao. Planning motion compliant to complex contact states. *Int. J. Robot. Res.*, 20(6):446–465, 2001.

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

[21] L. Kavraki, M. N. Kolountzakis, and J.-C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Trans. Robot. & Autom.*, 14(1):166–171, February 1998.

[22] L. Kavraki, J. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot motion planning. In *Proc. ACM Symp. on Theory of Computing*, pages 353–362, 1995.

[23] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. & Autom.*, 12(4):566–580, June 1996.

[24] Y. Koga, K. Kondo, J. Kuffner, and J. Latombe. Planning motions with intentions. *Computer Graphics (SIGGRAPH'94)*, pages 395–408, 1994.

[25] J. J. Kuffner and S. M. LaValle. RRT-connect: An efficient approach to single-query path planning. In *Proc. IEEE Int'l Conf. on Robotics and Automation*, 2000.

[26] F. Lamiraux and L. Kavraki. Planning paths for elastic objects under manipulation constraints. *Int. J. Robot. Res.*, 20(3):188–208, 2001.

[27] F. Lamiraux and L. Kavraki. Positioning and orienting a class of symmetric parts using a combination of a unit-radial and constant force fields: Computation of all equilibrium positions. *Int. J. Robot. Res.*, 20(8):635–659, 2001.

[28] J. Latombe. *Robot Motion Planning.* Kluwer, Boston, MA, 1991.

[29] S. LaValle and J. Kuffner. Randomized kinodynamic planning. *Int. J. Robot. Res.*, 5:348–400, May 2001.

[30] B. Mirtich and J. Canny. Impulse-based simulation of rigid bodies. In *Proceedings of ACM Interactive 3D Graphics Conference*, 1995.

[31] R. Motwani and P. Raghavan. *Randomized Algorithms.* Cambridge, 1996.

[32] J. R. Munkres. *Elements of Algebraic Topology.* Perseus Books, Cambridge Massachusetts, 1984.

[33] C. Nielsen and L. Kavraki. A two level fuzzy PRM for manipulation planning. In *IEEE/RSJ Int. Workshop on Intelligent Robots & Systems (IROS)*, pages 1716–1722, Japan, 2000.

[34] W. Rudin. *Principles of Mathematical Analysis.* McGraw-Hill, 1976.

[35] S. Sekhavat, P. Svestka, J.-P. Laumond, and M. H. Overmars. Multilevel path planning for nonholonomic robots using semiholonomic subsystems. *Int. J. Robot. Res.*, 17:840–857, 1998.

[36] P. Švestka. *Robot Motion Planning using Probabilistic Road Maps.* PhD thesis, Utrecht University, the Netherlands, 1997.

[37] P. Švestka and M. Overmars. Motion planning for car-like robots using a probabilistic learning approach. *Int. J. of Robotics Research*, 16(2):119–143, 1997.

[38] P. Svestka and M. H. Overmars. Coordinated motion planning for multiple car-like robots using probabilistic roadmaps. In *IEEE Int. Conf. Robot. & Autom.*, pages 1631–1636, 1995.

[39] R. Tombropoulos, J. Adler, and J. Latombe. Carabeamer: A treatment planner for a robotic radiosurgical system with general kinematics. *Medical Image Analysis*, 3(3):237–264, 1999.

[40] P. Valtr. On galleries with no bad points. Technical Report DIMACS Technical Report 96-55, DIMACS Center, Rutgers University, Piscataway, New Jersey, 1996.

Andrew M. Ladd received the B.S. degree in joint honors mathematics and computer science from McGill University, Montreal, Canada in 2000, and the M.S. degree in computer science from Rice University, Houston, TX in 2003. He is currently working toward the Ph.D. degree at Rice University. While at Rice he has worked on localization with wireless ethernet, analysis of PRM, parallel motion planning algorithms, knot untangling with motion planning, rope simulation for knotting, deformable models for fast image segmentation and is currently working on motion planning with dynamics.

Lydia Kavraki is an associate professor of Computer Science and Bioengineering at Rice University. Kavraki received her B.A. from the University of Crete in Greece and her Ph.D. from Stanford University in 1995. Her research interests are in robotics, and geometric computing, physical algorithms, bioinformatics and structural computational biology. A unifying theme in her work is the investigation of algorithms and system architectures for solving complex geometric problems arising in the physical world. She is the editor of one book and the author of more than seventy refereed publications. Kavraki has received the NSF CAREER award, an Alfred P. Sloan Research Fellowship, the Association for Computing Machinery Grace Murray Hopper Award for her work on probabilistic planning, the IEEE Robotics and Automation Society Early Career Award, and was selected as one of world's top 100 young innovators by the MIT Technology Review magazine in 2002.

IX. Appendix

A. Proofs of Measurability

The following provides proofs of Propositions III.1, III.2, and III.3 as well for Lemma III.4.

Proposition 1 (III.1) $f_n(x_1, ..., x_n)$ is a measurable function.

Proof: We begin by taking a decomposition of R into rectangles,

$$R = \bigcup_{i=1}^{\infty} A_i \times B_i.$$

We define $C_{ij} = A_i \cap B_j$ and consider the set

$$S^n = \bigcup_{i_1=1}^{\infty} \cdots \bigcup_{i_{n-1}=1}^{\infty} A_{i_1} \times C_{i_1 i_2} \times \cdots \times C_{i_{n-2} i_{n-1}} \times B_{i_{n-1}}.$$

We claim that $S^n = f_n^{-1}(1)$. $S^n \subset f_n^{-1}(1)$ is obvious by the construction of S^n. Suppose $f(x_1, ..., x_n) = 1$, then we have $i_1, ..., i_{n-1}$ such that $(x_{i_j}, x_{i_{j+1}}) \in A_{i_j} \times B_{i_j}$. This implies $x_{i_{j+1}} \in C_{i_j i_{j+1}}$ which demonstrates that $f_n^{-1}(1) \subset S^n$. By definition S^n is a measurable subset of X^n and we conclude that f_n is measurable. ∎

Proposition 2 (III.2) f_n^{xy} is measurable for every $x, y \in X$.

Proof: Let $x, y \in X$ be arbitrary and let $n \geq 0$ be given. Recall the definition of S^{n+2} and take a decomposition as before into rectangles $A_1^i \times \cdots \times A_{n+2}^i$. We claim that

$$(f_n^{xy})^{-1}(1) = \bigcup_{i:(x,y) \in A_1^i \times A_{n+2}^i} A_2^i \times \cdots \times A_{n+1}^i.$$

The claim is obviously true since $f_n(x, x_1, ..., x_n, y) = f_n^{xy}(x_1, ..., x_n)$. ∎

We will now show the transitive closure is measurable, which is Proposition III.3.

Proposition 3 (III.3) $\bar{R} \in \Sigma^2$.

Proof: Recall again the definition of S^n. We have a decomposition into disjoint rectangles for S^n,

$$S^n = \bigcup_{i=1}^{\infty} (A_i^n)_1 \times \cdots \times (A_i^n)_n.$$

We now claim that

$$\bar{R} = \bigcup_{n=2}^{\infty} \bigcup_{i=1}^{\infty} (A_i^n)_1 \times (A_i^n)_n.$$

This can be seen by observing that the definition

$$\bar{R} = \{(x, y) : \exists n, x_1, ..., x_n, f_{n+2}(x, x_1, ..., x_n, y) = 1\}$$

implies that $(x, y) \in \bar{R}$ if and only if there is n, i such that $(x, y) \in A_1^n \times A_n^n$. ∎

We now show that $\widehat{R}$ is measurable, which is Lemma III.4.

Lemma 1 (III.4) $\widehat{R}$ is measurable subset of X^2.

Proof: Following the proof in Proposition III.3, we begin by studying a decomposition of S^n into rectangles $(A_i^n)_j$. Let

$$\widehat{\mathcal{S}} := \left\{ (A_i^n)_1 \times (A_i^n)_n : n \geq 2, \mu^{n-2} \left(\prod_{j=2}^{n-1} (A_i^n)_j \right) > 0 \right\},$$

and then claiming that $\bigcup \widehat{\mathcal{S}} = \widehat{R}$. The left hand inclusion of $\widehat{\mathcal{S}} \subseteq \widehat{R}$ is easily seen. To show the reverse inclusion, suppose Algorithm 2 succeeds with probability greater than 0 on some query (x, y) but $(x, y) \notin \bigcup \widehat{\mathcal{S}}$. We observe that for some $n \geq 0$, there is probability greater than 0 that on iteration n, the algorithm reaches that iteration and succeeds. The case $n = 0$ is trivial so we do not consider it. We call this probability p and observe that it is equal to the probability of choosing $x_1, ..., x_n$ at random from X^n such that $f_i^{xy}(x_1, ..., x_n) = 1$. Observe that

$$p = \mu^n \left((f_n^{xy})^{-1}(1) \right).$$

By the definition of f_n^{xy} used in Proposition III.2 and since the inequality

$$0 < p = \mu^n((f_n^{xy})^{-1}(1)) \leq \sum_{i:(x,y) \in (A_i^{n+2})_1 \times (A_i^{n+2})_{n+2}} \mu^n \left(\prod_{j=2}^{n-1} (A_i^{n+2})_j \right),$$

implies there exists i such that

$$\mu^n \left(\prod_{j=2}^{n+1} (A_i^{n+2})_j \right) > 0.$$

It follows that $(A_i^{n+2})_1 \times (A_i^{n+2})_{n+2} \in \widehat{\mathcal{S}}$ and $(x, y) \in \bigcup \widehat{\mathcal{S}}$ which contradicts the initial assumption. We conclude that $\widehat{R} = \widehat{\mathcal{S}}$, a measurable set by construction. ∎

B. CW Complex Proof

In this subsection, we give a proof for Theorem VI.1.

Theorem 1 (VI.1) Let X be a finite CW complex with decomposition $X_0 \subset \cdots \subset X_n$. Suppose we are given a measure μ and a local planner R such that for every cell B,
1. if $A \subset B$ is open and non-empty in B, then $\mu(A) > 0$,
2. if xRy with $y \in Int(B)$, there there exists $A \subset Int(B)$ such that A is non-empty and open in B and $y' \in A$ implies xRy',
3. R is symmetric and $\bar{R}$ contains path reachability for B.
It follows that a PRM planner using μ and R is probabilistically complete for computing path reachability in X.

Proof: This proof requires that we show that the conditions of Theorem IV.2 are met. For each individual cell, it is very easy to check that the conditions for Theorem IV.2 hold using the measure and planner restricted to that cell. We will explain why it also holds for the compound space.

Let $x, y \in X$ be any two path connected points and let $\gamma : [0,1] \to X$ be the path that connects them. Suppose there exists i such that $\gamma^{-1}(B_i)$ is disconnected. This implies there is $a < b \in [0,1]$ that lie in different components of $\gamma^{-1}(B_i)$. Since $\gamma(a)$ and $\gamma(b)$ are path connected in B_i, there exists a path from x to y, γ', such that for $0 \le c \le a$ and $b \le c \le 1$, $\gamma'(c) = \gamma(c)$ and for $a \le c \le b$, γ' stays in B_i. Therefore without loss of generality, the preimage $\gamma^{-1}(B_i)$ is connected for all i.

Using the preimages, $\gamma^{-1}(B_i)$, we obtain a partition of the interval $[0,1]$ into blocks:

$$[a_0, a_1], ..., [a_i, a_{i+1}], ..., [a_{M-1}, a_M]$$

where $a_0 = 0$, $a_M = 1$. For each point $c \in [a_i, a_{i+1}]$ in a block, there is a cell B_i such that $\gamma(c) \in B_i$. For each $0 \le i < j \le M$, observe that $B_i \ne B_j$. For each i, $\gamma(a_i)$ and $\gamma(a_{i+1})$ are both in some B_i. Since each B_i is path connected, there is a sequence of points in B_i, $b_1, ..., b_k$ such that $\gamma(a_i) R b_1 R ... R b_k R \gamma(a_{i+1})$. So over each block there is a local planner path through the corresponding ball. We can now paste these paths together. We conclude that there exists $x_1, ..., x_m \in X$ such that $xRx_1R...Rx_mRy$.

Having a shown that a path using R exists between x and y exists if they are path connected, we must show that the path satisfies the conditions for Theorem IV.2. To do this, it is sufficient to prove that for any three points $x_a, x_b, x_c \in X$ where $x_a R x_b R x_c$ that there exists a measurable subset A of X such that for all $z \in A$, $x_a R z R x_c$ and $\mu(A) > 0$.

Observe that there is a unique B such that $x_b \in Int(B)$. By assumption, there is A_1 an open subset in B such that $z \in A_1$ implies $x_a R z$. Similarily, there is A_2 an open subset in B such that $z \in A_2$ implies $z R x_c$. Now $A = A_1 \cap A_2$ and $x_b \in A$ implies $\mu(A) > 0$ since A is open in B. ■

C. Polygon Robot with Contact

In this subsection, we provide proofs for the lemmas and the theorem of Subsection VI-C. We begin by proving Lemma VI.3 which follows from the given definitions.

Lemma 2 (VI.3) The Z_α are pairwise disjoint and their union is X.

Proof: Suppose there was α, β such that $\alpha \ne \beta$ and $Z_\alpha \cap Z_\beta \ne \emptyset$. Let $\gamma = \alpha \cup \beta$ and observe $Z_\alpha \cap Z_\beta \subset X_\gamma$. Now $\alpha \prec \gamma$ since $\alpha \ne \beta$ and therefore $x \in Z_\alpha \cap Z_\beta$ implies that $\gamma(x)$. This contradicts the definition of Z_α and therefore no such x exists.

Let $x \in X$ be arbitrary. Suppose $\alpha(x)$ and notice that $x \notin Z_\alpha$ implies that there is β with $\alpha \prec \beta$ such that $\beta(x)$. We proceed recursively by noting that $\emptyset(x)$ and that this sequence eventually terminates as the set of all constraints is the largest set in the partial order. Therefore, the union of Z_α is X as every point of x appears in some unique Z_α. ■

To proceed with the next proof, we need to define parametrizations for the equivalence classes of the contact constraints. Throughout, this section we will assume certain geometric degeneracies do not occur. Precisely, no s_i, s_j parallel and no p_i, p_j, p_k collinear. The same is assumed for obstacle vertices and edges. The proofs we provide can be amended if the general position assumption is removed, however it greatly complicates the presentation and is beyond the scope of this paper. We note that the Algorithms 3 and 4 are still correct without the general position assumption although it requires additional checking.

To each $\alpha \in \mathcal{P}$, we can assign a subspace of X which we call X_α. X_α is the subspace in which the constraint pairs in X are realized. More formally,

$$X_\alpha = \{x \in X : \alpha(x)\}.$$

For example, if $\alpha = \{(v_i, s_j)\}$, then X_α is the subspace of X consisting of configurations where obstacle vertex v_i intersects with some point of s_j.

The collection $\mathcal{P}$ forms a partial order ordered by the relation $\preceq$. This induces a partial order over the subspaces. In particular, if $\alpha \subset \beta$ then X_β is a subspace of X_α. Also, note that $X = X_\emptyset$.

The general scheme will be that for every α, there will be an integer $0 \le k_\alpha \le 3$, a k_α-manifold with boundary A_α and an injective, continuous map

$$\pi_\alpha : A_\alpha \to T.$$

We will further prove that $\pi^{-1}(X)$ is homeomorphic to X_α. Once the parametrizations are in hand, we will be ready to prove Lemma VI.4.

Let α be some constraint set. We will now enumerate various cases for α:
1. $\emptyset$.
2. $\{(p_i, e_j)\}$.
3. $\{(p_i, e_j), (p_i, e_k)\}$ where $e_j \cap e_k = v_l$.
4. $\{(p_i, e_j), (p_k, e_l)\}$ where $i \ne j$.
5. $\{(s_i, v_j), (p_k, e_l)\}$.

We note that the above list is not exhaustive. Each of the non-empty cases has a dual obtained by treating the polygon P as fixed and allowing the obstacles O to move. For example, a constraint set of type $\{(s_i, v_j)\}$ is dual to a constraint set of type $\{(p_i, e_j)\}$. We now examine each constraint set above separately and define in each case A_α and π_α. In order to do that we need to define two geometric maps. Let $\psi : P \times S^1 \times \mathcal{R}^2 \to T$ be the transformation given by the rule $\psi(p, \theta, x)$ rotates P by θ about point $p \in P$ and then translates by x to obtain an embedding in T. If s is any line segment then $s : [0,1] \to \mathcal{R}^2$ is the parametrization for the segment.

We continue now by describing the parametrizations. The cases can be easily verified by geometric arguments.

Case 1: Suppose $\alpha = \emptyset$ then set $A_\alpha = T$ and say that π_α be the identity map on T. $\pi_\alpha^{-1}(X)$ is homeomorphic to $X_\alpha = X$.

Case 2: Suppose $\alpha = \{(p_i, e_j)\}$. We set $A_\alpha = [0,1] \times S^1$ and use

$$\pi_\alpha(c, \theta) = \psi(p_i, \theta, e_j(c) - p_i).$$

π_α is an injection and $\pi_\alpha^{-1}(X)$ is homeomorphic to X_α.

Case 3: Suppose $\alpha = \{(p_i, e_j), (p_i, e_l)\}$ and so that $e_j \cap e_l = v_m$. We set $A_\alpha = S^1$ and use

$$\pi_\alpha(\theta) = \psi(p_i, \theta, v_m - p_i).$$

π_α is an injection and $\pi_\alpha^{-1}(X)$ is homeomorphic to X_α.

Case 4: Suppose $\alpha = \{(p_i, e_j), (p_k, e_l)\}$ where $i \neq k$. We set

$$A_\alpha = \{(\theta, c) : \psi(p_i, \theta, e_j(c) - p_i)(p_k) \in e_l\},$$

and use the map

$$\pi_\alpha(\theta, c) = \psi(p_i, \theta, e_j(c) - p_i).$$

π_α is an injection and $\pi_\alpha^{-1}(X)$ is homeomorphic to X_α. Also for every $\theta \in S^1$ there is at most one $c \in [0,1]$ such that $(\theta, c) \in A_\alpha$.

Case 5: Suppose $\alpha = \{(s_i, v_j), (p_k, e_l)\}$. We set

$$A_\alpha = \{(\theta, c) : \psi(s_i(0), \theta, v_j - s_i(c))(p_k) \in e_l\},$$

and

$$\pi_\alpha(\theta, c) = \psi(s_i(0), \theta, v_j - s_i(c)).$$

π_α is an injection and $\pi_\alpha^{-1}(X)$ is homeomorphic to X_α. Also for every $\theta \in S^1$ there is at most one $c \in [0,1]$ such that $(\theta, c) \in A_\alpha$.

Finally, we argue that any other constraint type is either unsatisfiable or consists of 0, 1 or 2 points. This argument makes heavy use of the general position assumption. There are additional exceptions if this condition is removed.

Proposition IX.1: Let α be any constraint that does not appear in the list or as a dual of a type in the list. Then X_α consists of either 0, 1 or 2 points and hence A_α, π_α are trivial.

Proof: If $\{(p_i, e_j), (p_i, e_k)\} \subset \alpha$ where $e_j \cap e_k = \emptyset$ then $X_\alpha = \emptyset$ since a point cannot lie on two disjoint lines. Otherwise $|\alpha| \geq 3$. Suppose $|\alpha| = 3$. By taking each of the three constraint types with two elements and adding a (p_i, e_j) constraint, we can obtain all types with 3 elements up to duality. Consider the set

$$C = \{\pi_\alpha(a)(p_i) : a \in A_\alpha\}.$$

Observe that C is a curve and that it lies on the boundary of its convex hull. Additionally, it is important to note that the general position assumption insures that C is not a straight line and is smooth. Thus, any line segment intersects it at most twice. In particular, e_j intersects it at most twice, demonstrating that X_α consists of 0, 1 or 2 points. Now if $|\alpha| > 3$, we know that that there $\beta \prec \alpha$ with $|\beta| = 3$ and $X_\alpha \subset X_\beta$ finishes the result. It is important to note that X_α may not be empty when $|\alpha| > 3$ even with the general position assumption. ■

We turn now to the spaces Z_α and show that they are open manifolds.

Lemma 3 (VI.4) For every α, either $Z_\alpha = \emptyset$ or there is an integer $0 \leq k_\alpha \leq 3$ such that Z_α is a k_α-manifold.

Proof: For some α, suppose Z_α is not empty. We claim then that for Case 1 that $k_\alpha = 3$, for Case 2 that $k_\alpha = 2$, for Case 3, Case 4 and Case 5 that $k_\alpha = 1$ and for any other situation that $k_\alpha = 0$. This argument is based on counting the number of free variables in A_α. We point out that this reasoning is sound since a degenerate A_α can only arise when Z_α is empty.

Let $x \in Z_\alpha$ and observe that by definition on the constraints in α are satisfied at configuration x. It follows that there exists $\epsilon > 0$ such that for every $y \in B_\epsilon(x)$ such that $\alpha(y)$, there is no β with $\alpha \prec \beta$ and $\beta(y)$. Also, observe that $C = \pi_\alpha^{-1}(B_\epsilon(x))$ is open in A_α. Finally, $\pi_\alpha(C) \subset X$ by construction of the neighbourhood around x. We conclude that every point $x \in Z_\alpha$ has neighbourhood C_x which is open in A_α. Each neighbourhood has dimension k_α proving that Z_α is a k_α-manifold. ■

We are finally ready to prove the main theorem for this section.

Theorem 2 (VI.2) X is a finite CW complex. There is an explicit measure and local planner for X that satisfy the conditions of Theorem VI.1.

Proof: A technicality which complicates this proof is that Z_α which $k_\alpha \geq 2$ might have holes. For the moment, we will assume this cannot happen however by relaxing the definition of CW complex to allow the cells to be manifolds with boundary rather than closed balls we can obtain a slight variation of Theorem VI.1 which has essentially the same proof. The argument for X being a CW complex also holds under the assumption that there are no holes.

We submit that $\overline{Z_\alpha}$ are the cells and that

$$X_d = \bigcup_{\alpha : k_\alpha \leq d} Z_\alpha.$$

The X_d are the skeletons for the CW complex structure.

Lemma VI.4 proves that $\overline{Z_\alpha}$ are closed k_α-manifolds for $0 \leq k_\alpha \leq 3$ or are empty. The cell structure is correct for $d = 0$ and for $d \geq 0$, let α be such that $k_\alpha = d + 1$. In each case, we observe $Bd(Z_\alpha) \subset X_d$ and then take f to be the identity on X restricted to $Bd(Z_\alpha)$. This is continuous by definition. Because the cells are pairwise disjoint, it follows that this is a correct inductive definition of a CW structure. This follows from Lemma VI.3 as does the fact that $X_3 = X$ when together with the observation in Proposition IX.1.

Since there are only finitely many Z_α, we can conclude that it is a finite CW complex.

Now for any α with Z_α non-empty, let $A \subset \overline{Z_\alpha}$ be open and non-empty in $\overline{Z_\alpha}$. Let p be the probability of choosing α at random in Algorithm 3 and observe that $p > 0$. Let μ_α be the uniform measure for A_α and observe that Algorithm 3 uses this measure to induce a measure for sampling a point from X satisfying α conditional on α being selected. It follows that measure μ induced by Algorithm 3 satisfies

$$\mu(A) \geq p\mu_\alpha(A) > 0,$$

as desired.

Suppose for x, y that xRy. Now let α be maximal such that $\alpha(x)$ and $\alpha(y)$. Without loss of generality, we can infer that $x \in \overline{Z_\alpha}$ and $y \in \overline{Z_\alpha}$. By appealing to results such as those presented in [5] for manifolds with or without boundary, we observe that there is A, open in $\overline{Z_\alpha}$ such that $z \in A$ implies xRz and A is neighbourhood of y. Now let β be such that $y \in Z_\beta$. This is unique because of Lemma VI.3. $A' = A \cap Z_\beta$ is open in Z_β and non-empty because $y \in A'$.

Finally, we note that R is clearly symmetric and that $\bar{R}$ contains path reachability for every Z_α. The final point can be argued by the path isolation argument for manifolds presented in [5] and in other works. ■

The theorem still holds without the general position assumption, however it requires taking more cases for the parametrizations because of geometric degeneracies that can occur. The proof is beyond the scope of this paper.