2013 IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS)
November 3-7, 2013. Tokyo, Japan

A Study on the Finite-Time Near-Optimality Properties of Sampling-Based Motion Planners

Andrew Dobson, Kostas E. Bekris

***Abstract*— Sampling-based algorithms have proven practical in solving motion planning challenges in relatively high-dimensional instances in geometrically complex workspaces. Early work focused on quickly returning feasible solutions. Only recently was it shown under which conditions these algorithms asymptotically return optimal or near-optimal solutions. These methods yield desired properties only in an asymptotic fashion, i.e., the properties are attained after infinite computation time. This work studies the finite-time properties of sampling-based planners in terms of path quality. The focus is on roadmap-based methods, due to their simplicity. This work illustrates that existing sampling-based planners which construct roadmaps in an asymptotically (near-)optimal manner exhibit a "probably near-optimal" property in finite time. This means that it is possible to compute a confidence value, i.e. a probability, regarding the existence of upper bounds for the length of the path returned by the roadmap as a function of the number of configuration space samples. This property can result in useful tools for determining existence of solutions and a probabilistic stopping criterion for PRM-like methods. These properties are validated through experimental trials.**

I. Introduction

Traditional sampling-based motion planning approaches, such as PRM [9] and RRT [12], efficiently and quickly provide feasible solutions, even for relatively high-dimensional problem instances. These methods are probabilistically complete [4], [8], [11]. Recent progress in the field has also provided the conditions under which these methods converge to optimal solutions [7]. These properties hold asymptotically, i.e., after an infinite amount of computation time. This implies that when a sampling-based motion planner is practically terminated after a finite amount of time given a typically ad-hoc stopping criterion, there is no guarantee for the quality of the path returned.

This work seeks to identify the finite-time properties of sampling-based motion planners, so as to address this limitation. The reasoning in this paper follows a direction in previous work for analyzing the properties of sampling-based planners [7], [8]. A set of covering hyper-balls over optimal paths can be defined as shown in Figure 1. There is positive probability of generating samples within these balls, and it is also possible to pairwise connect these samples through the collision-free configuration space. This results in a worst-case upper bound for a variant of PRM*.

Work by the authors has been supported by NSF CNS 0932423. Any conclusions expressed here are of the authors and do not reflect the views of the sponsors. The authors are with the Computer Science Department, Rutgers University, 110 Frelinghuysen Road, 08854, Piscataway, NJ, USA, email: kostas.bekris at cs.rutgers.edu.

The covering ball upper bound on path length is not very useful, as it does not change as the number of samples drawn increases. This is due to the fact that even though the radius of the balls decreases over time, their number increases. In practice, it is well known that a greater number of samples reduces path length. To understand this behavior, this work employs a set of disjoint balls, as illustrated in Figure 1, of a fixed number, M, but of decreasing radius as a function of the number of samples. Bounds can then be drawn for the expected length of a path through the disjoint balls, and a probabilistic stopping criterion to traditional PRM-like algorithms is defined. This analysis results in the introduction of a new, proposed property for sampling-based motion planners, that of *probabilistic near-optimality* (PNO). This is a similar type of guarantee provided by survey sampling, where there is a bound regarding the error of a poll with a confidence value that the result of the survey will remain within this error.

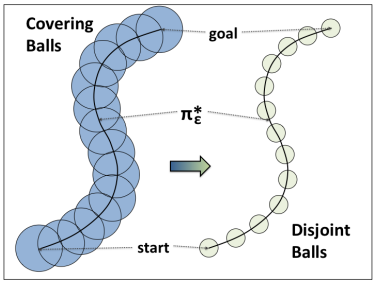

Fig. 1. "Disjoint Balls" are concentric with their respective "Covering Balls" with radius r, which is a diminishing function of the number of samples, n.

This work takes a first step in analyzing the finite-time properties of sampling-based motion planners and introduces new interesting questions which warrant further exploration. A model is drawn and is tested through Monte Carlo experiments, and it is shown that models for probability of path coverage and expected path length hold well for basic systems and environments.

II. Related Work

An early technique which popularized sampling-based planning is the Probabilistic Roadmap Method (PRM) [9]. The method pre-processes a configuration space C, and produces a roadmap, i.e., a graph in C, which can answer multiple queries. The method is probabilistically complete in general setups [4], [8], [11], and there are efforts in showing solution non-existence [15], [19]. There are also extensions of the PRM, which attempt to improve path quality, such as covering different homotopic classes [5], [16]. After the success of PRM, tree-based analogs, such as RRT [12], were also proposed.

With the proposal of new sampling-based planners [6], [7], it was shown that variants of sampling-based planners,

such as RRT* and PRM*, are asymptotically optimal. These methods build the most sparse planning structures possible to achieve optimality, but the resulting graphs are still relatively dense. If optimality is relaxed, it is possible to construct even sparser structures by taking advantage of work on graph spanners [14]. Graph spanners are subgraphs of their input graphs with guarantees on path costs. This results in asymptotic near-optimality. An incremental version provides good path quality [13] and it is possible to achieve such properties without including all configurations as nodes in the resulting graph [2], [3].

This work focuses on *probabilistic near-optimality*. This property is inspired by a similar guarantee of *probably approximately correct* (PAC) algorithms [18]. PAC algorithms yield a probabilistic confidence bound on finding a solution. The connection between sampling-based machine learning techniques and sampling-based motion planning has been examined before [10]. This work seeks to identify the *probably near-optimal* (PNO) properties of sampling-based planners in geometrically complex configuration spaces.

III. Setup

This section provides definitions for understanding the proposed property and the corresponding analysis.

Definition *1 (Configuration Space):* A configuration space, denoted C, contains as point elements the configurations of a moving system. This space is d-dimensional, where each of the d axes typically corresponds to a degree of freedom of the robot.

C can be partitioned into two sets; the valid (C_{free}) and invalid (C_{obs}) system configurations. C is also endowed with a metric distance function.

Definition *2 (Distance Function):* Given two points in C: q_a and q_b, the function $dist(q_a, q_b)$ is a metric function that returns a real value, i.e., $dist(q_a, q_b) : (C \times C) \to \mathbb{R}$.

The objective of sampling-based methods is to produce collision-free paths, which travel between points in C.

Definition *3 (Path):* A path is a continuous function π, defined over the domain $[0, 1]$ and range C.

Then, the path planning problem is defined as follows.

Definition *4 (The Path Planning Problem):* Given C_{free} and two points within this space, q_{start}, q_{goal}, representing an initial and a goal configuration, find a continuous path $\pi : [0, 1] \to C_{\text{free}}$, such that $\pi(0) = q_{start}$ and $\pi(1) = q_{goal}$.

Paths through C_{free} may pass infinitely close to the invalid set C_{obs}. A path's clearance is the minimum distance from a point in the path to C_{obs}.

Definition *5 (Robust Feasible Path):* A path, π_ϵ, is ϵ-robust if π_ϵ has clearance of at least ϵ.

Sampling-based planners employ basic primitives, one of which is a local planner, which in this work produces a straight-line path, denoted $L(q_a, q_b)$ where q_a, $q_b \in C_{\text{free}}$. The volume of a set, S, is denoted $|S|$, where this work employs the Lebesgue measure.

A. Properties of Sampling-Based Algorithms

It is well-known that sampling-based planners provide probabilistic completeness [8].

Definition *6 (Probabilistic Completeness):* A probabilistically complete algorithm returns a solution with probability 1 if such a solution exists, given infinite computation. If no solution has been returned, it is unknown if a solution exists.

Another relevant property is asymptotic optimality [6], [7].

Definition *7 (Asymptotic Optimality):* As the iterations n of an asymptotically optimal algorithm increase, the algorithm converges to returning π^*_ϵ.

Asymptotically optimal planners, such as RRT* and PRM* result in relatively dense data structures. An extension to alleviate this problem are planners which provide asymptotic near-optimality [3], [13].

Definition *8 (Asymptotic Near-Optimality):* As the iterations of an asymptotically near-optimal algorithm go to infinity, the algorithm returns a solution converging to path length $|\pi_{\texttt{PNO}}|$ upper bounded by $\leq a \cdot |\pi^*_\epsilon| + b$ for $a, b \in \mathbb{R}^+$.

These methods compute compact representations which are quick to query, and have low impact on path cost. This work shows that asymptotically (near-)optimal planners exhibit a similar property in finite time.

Definition *9 (Probabilistic Near-Optimality):* At a finite iteration n, a probably near-optimal algorithm provides a confidence value, $0 \leq p < 1$, representing the probability that it can provide a solution $\pi_{\texttt{PNO}}$ that can be upped bounded relative to the optimum one, i.e., such that $|\pi_{\texttt{PNO}}| \leq a \cdot |\pi^*_\epsilon| + b$ for known constants $a, b \in \mathbb{R}^+$.

B. Description of Specific Algorithmic Variant

Algorithm 1: PNO-PRM*(n, ϵ)

$V \leftarrow \{\texttt{SampleFree}_i\}_{i=1,\ldots,n}$; $E \leftarrow \emptyset$;
for $v \in V$ **do**
  $r_m \leftarrow max\{\frac{3\epsilon}{2}, \gamma_{PRM}(log(n)/n)^{1/d}\}$;
  $U \leftarrow \texttt{NEAR}(V, v, r_m)$;
  for $u \in U$ **do**
    if $L(v, u) \in C_{free}$ **then**
      $E \leftarrow E \cup \{L(v, u)\}$;
return $G = (V, E)$;

This work considers a PRM* variant dubbed PNO-PRM*, shown in Algorithm 1. The algorithm takes as parameters n, the number of samples to generate, and ϵ, the desired clearance. The value for n is based on a probabilistic stopping criterion, detailed in Section IV.F. The algorithm also imposes a minimum connection distance of $\frac{3\epsilon}{2}$ for pairs of samples. This connection distance comes as a result from the analysis (Corollary 2).

IV. Analysis

The analysis begins with some preliminary material. Following these, individual algorithms are examined in greater detail where bounds on path cost in finite time are discussed.

A. Probability of Path Coverage

This work follows the terminology used in prior work [7], [8], which defines a volume around an arbitrary path, π, via a set of covering balls.

Definition *10 (Covering Balls):* Given a path, $\pi : [0,1] \rightarrow C$, define a finite set of covering balls as $M = \lceil \frac{2 \cdot |\pi_\epsilon^*|}{\epsilon} \rceil + 1$ hyper balls, $\{B_0(\frac{\epsilon}{2}), B_1(\frac{\epsilon}{2}), ..., B_{M-1}(\frac{\epsilon}{2})\}$, of radius $\frac{\epsilon}{2}$ centered along π and satisfying:

- The first ball, $B_0(\frac{\epsilon}{2})$, is centered at $\pi(0)$, and the last ball $B_{M-1}(\frac{\epsilon}{2})$ is centered at $\pi(1)$.
- The centers of two consecutive balls are at most $\frac{\epsilon}{2}$ apart, and the distance between all consecutive pairs is uniform.

Assume there exists a robustly optimal path $\pi_\epsilon^* \subset C_{\text{free}}$, which answers a specific query. The clearance, ϵ, must be sufficiently small in order for this path to travel through narrow passages. Figure 2 gives an illustration of a construction of covering balls over π_ϵ^*. Consider the probability of sampling in each of these balls.

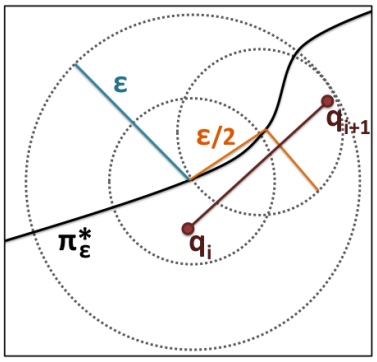

Fig. 2. Covering balls over an optimal path with radius and separation $\frac{\epsilon}{2}$. All paths $L(q_i, q_{i+1})$ are in the valid part of the space.

Lemma *1 (Probability of Path Coverage):* For some arbitrary ϵ-robust optimal path, π_ϵ^*, the probability of having at least one sample in each covering ball is

$$P_{cover} = \left(1 - \left(1 - \frac{|B(\frac{\epsilon}{2})|}{|C_{\text{free}}|}\right)^n\right)^M$$

where n is the number of samples produced.

Proof This proof follows a similar argument to that provided by Kavraki and Kolountzakis [8] but draws a tighter bound on this probability. If the algorithm generates n samples, the probability that no sample lies in the ball is a set of n Bernoulli Trials. The probability with which a sample falls into a ball is $\frac{|B(\frac{\epsilon}{2})|}{|C_{\text{free}}|}$, resulting in a probability of $(1 - \frac{|B(\frac{\epsilon}{2})|}{|C_{\text{free}}|})^n$, to cover this ball. The overall probability then is another set of Bernoulli Trials, assuming sample independence. There are at most $M = \lceil \frac{2 \cdot |\pi_\epsilon^*|}{\epsilon} \rceil + 1$ balls along this path. Taking the compliment of the probability of having zero samples in one ball gives the probability of having one or more samples in this ball, which is $(1 - (1 - \frac{|B(\frac{\epsilon}{2})|}{|C_{\text{free}}|})^n)$. M such trials yields the proposed probability. □

Figure 3 provides a Monte Carlo simulation in a corridor environment, which tests the validity of P_{cover}. After covering π_ϵ^*, the algorithm must connect samples pair-wise in consecutive balls. Essentially, the algorithm must test for connections up to the maximum possible distance between samples.

Lemma *2 (Distance between Covering Samples):* Given two arbitrary points within consecutive covering balls, $q_n \in B_n$ and $q_{n+1} \in B_{n+1}$, then $dist(q_n, q_{n+1}) \leq \frac{3\epsilon}{2}$.

Proof Each ball has radius $\frac{\epsilon}{2}$, thus, the distance from samples q_n and q_{n+1} can be no more than $\frac{\epsilon}{2}$ from the center of their respective balls, B_n and B_{n+1}. Furthermore,

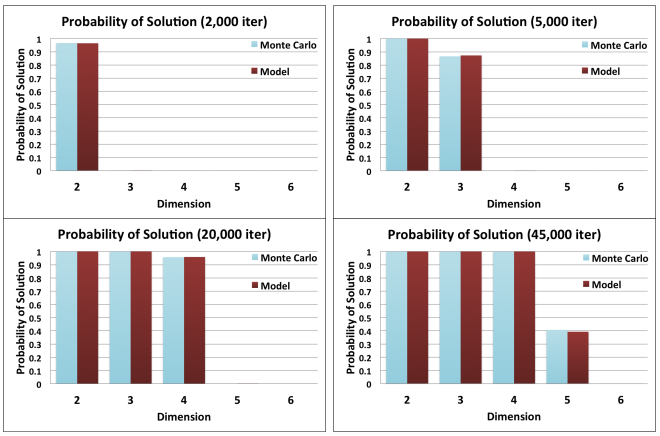

Fig. 3. Probability of answering a query in the illustrated environment according to the model against Monte Carlo experiments for varying numbers of samples. For the different dimensions of the space, x_i, x_1 is bounded by $[0, 100]$ while all other x_i are bounded by $[0, 2]$. $|\pi_\epsilon^*| = 94$, $M = 95$, $r = 0.5$. Monte Carlo values are averaged over 1000 trials.

by construction, the centers of these balls are no more than $\frac{\epsilon}{2}$ apart. It follows then that the distance between q_n and q_{n+1} can be no more than $\frac{3\epsilon}{2}$. □

The paths between the samples therefore lie entirely in the free space C_{free}.

Lemma *3 (Connectivity of Covering Samples):* Given two points within consecutive covering balls, $q_n \in B_n$ and $q_{n+1} \in B_{n+1}$, then $L(q_n, q_{n+1})$ is collision-free.

Proof Consider two consecutive covering balls, B_n and B_{n+1}, as well as the hyper ball centered around the center of B_n with radius ϵ as in Figure 2, and call this ball E. It follows intuitively that $B_n \subset E$. It is also the case that $B_{n+1} \subset E$, as the distance from the center of E to the center of B_{n+1} is no more than $\frac{\epsilon}{2}$ and B_{n+1} also has radius $\frac{\epsilon}{2}$. It follows then that for any arbitrary pair of points, $q_n \in B_n$ and $q_{n+1} \in B_{n+1}$, that $L(q_n, q_{n+1}) \subset E$. By definition, $E \subset C_{\text{free}}$, thus, $L(q_n, q_{n+1}) \subset C_{\text{free}}$. □

Combined, lemmas 1 - 3 provide the probability of successfully answering a query with a path, which covers a specific π_ϵ^* and has length at most $(3 \cdot |\pi_\epsilon^*| + \epsilon)$.

B. Determining Path Length

Analyzing path cost using the above lemmas produces a length bound invariant of the number of samples in the space. Any path in the roadmap will be no longer than 3 times the robustly feasible optimum path. This is because the longest possible paths between samples in neighboring covering balls are 3 times the displacement between the balls' centers ($\frac{\epsilon}{2}$). Decreasing the displacement between covering balls also results in a higher number of balls to cover a path, resulting in constant worst-case bound. In order to create bounds which depend on n, a set of disjoint balls is defined.

Definition *11 (Disjoint Balls):* Given a path π_ϵ^*, define the disjoint balls as the set of $M = \lceil \frac{2 \cdot |\pi_\epsilon^*|}{\epsilon} \rceil + 1$ balls, $\{B_0(r), B_1(r), ...B_{M-1}(r)\}$, with radius:

$$r = \left(\frac{(\frac{d}{2})! \cdot |C_{\text{free}}| \cdot (1 - (1 - p^{\frac{1}{M}})^{\frac{1}{n}})}{\pi^{\frac{d}{2}}} \right)^{\frac{1}{d}}$$

with $B_0(r)$ centered at $\pi(0)$, $B_{M-1}(r)$ centered at $\pi(1)$ and the distance between all consecutive balls is uniform.

The relation between covering balls and disjoint balls is shown in Figure 1 (Page 1). The analysis discusses the situation after sufficient samples have been drawn to ensure r is small enough to guarantee that disjoint balls are in fact disjoint. The prior lemmas extend to the case of disjoint balls.

Corollary *1 (Probability of Disjoint Coverage):* For n samples, the probability that there is a sample in each of the disjoint balls over π^*_ϵ is

$$P_{cover} = \left(1 - \left(1 - \frac{|B(r)|}{|C_{\text{free}}|}\right)^n\right)^M.$$

Corollary *2 (Distance between Disjoint Samples):* Given two arbitrary points within consecutive disjoint balls, $q_n \in B_n$ and $q_{n+1} \in B_{n+1}$, then $dist(q_n, q_{n+1}) \leq \frac{\epsilon}{2} + 2r$.

Corollary *3 (Connectivity of Disjoint Samples):* Given two points within consecutive disjoint balls, $q_n \in B_n$ and $q_{n+1} \in B_{n+1}$, then $L(q_n, q_{n+1})$ is collision-free.

Again, if a sampling-based algorithm attempts all connections for samples within the distance in Corollary 2, then the probability of covering π^*_ϵ is P_{cover}. To begin reasoning about the length of constructed paths, a δ value is used.

Definition *12 (δ-value):* Consider two consecutive disjoint balls, B_n, B_{n+1}, centered at c_n and c_{n+1} respectively. Two samples q_n, q_{n+1} are drawn uniformly within these balls. Then, δ is the ratio of the expected distance between the samples to the distance between the ball centers, i.e., $\delta = \frac{E[dist(q_n, q_{n+1})]}{dist(c_n, c_{n+1})}$.

Calculating δ is a difficult problem, related to the "ball line picking" problem [17], for which an analytical solution is not obvious. Nevertheless, Monte Carlo experiments provided in Figure 4 indicate this value in $\mathbb{R}^d$. Note that these values are generated using two samples in unit-displaced hyper-spheres in obstacle-free $\mathbb{R}^d$-space.

For the remaining discussion, the paper will refer to "returned", "average", and "worst"-case paths. A "returned" path refers to the path returned by an algorithm. An "average" path refers to any path, which answers the query, and whose vertices all lie within the disjoint balls. A "worst"-case path refers to a hypothetical path, which contains a single transition vertex in each disjoint ball covering the optimum path that has maximum length.

C. Finite-Time Properties of PNO-PRM*

The discussion on the provided bounds in path quality begins with the PNO-PRM* algorithm.

1) Worst-Path Analysis: The worst path constructed through disjoint balls is bounded in length.

Theorem *1 (Worst Case Path Degradation):* For an arbitrary query, with probability P_{cover}, PNO-PRM* yields an answer of length $\ell \leq (1 + 2r) \cdot |\pi^*_\epsilon| + 2r$.

Proof Consecutive disjoint balls have separation at most $\frac{\epsilon}{2}$ and radius r. Samples in these balls have maximum distance $(\frac{\epsilon}{2} + 2r)$. Each of the segments along $|\pi^*_\epsilon|$ is covered with length at most $(\frac{\epsilon}{2} + 2r)$. This yields a total length of $(\frac{\epsilon}{2} + 2r)(M - 1) = (1 + 2r) \cdot |\pi^*_\epsilon|$. Adding the maximum $2r$ additive connection cost yields the above bound. □

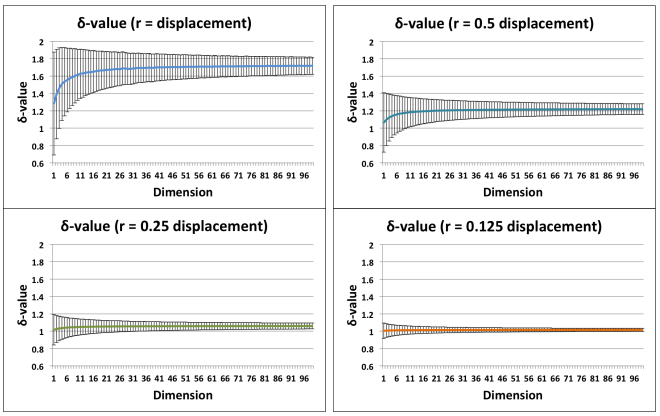

Fig. 4. Monte Carlo simulation data for δ in various dimensions. Individual graphs show different 'ball radius'-to-'ball displacement' ratios. Data was obtained by sampling in dislocated balls and calculating sample distance.

2) Average-Path Analysis: This analysis employs the δ value described in Definition 12. This δ value ends up representing a multiplicative bound on path length.

Theorem *2 (Average Path Degradation):* Given the average relative distance between two points in consecutive disjoint balls, δ, along some arbitrary path answering an arbitrary query, with probability P_{cover} (Lemma 1), the graph of PNO-PRM* contains a path with expected length:

$$E[\ell] = \delta \cdot |\pi^*_\epsilon| + 2(\frac{d \cdot r}{d+1})$$

Proof An illustration of an average path over a query is shown in Figure 5. The expected returned path length for each segment is the length of this segment times δ. By construction (Definition 10), there are at most $M = \lceil \frac{2 \cdot |\pi^*_\epsilon|}{\epsilon} \rceil + 1$ such balls, which yields $\lceil \frac{2 \cdot |\pi^*_\epsilon|}{\epsilon} \rceil$ segments to traverse, each at relative cost δ. This results in a total cost of $\delta \cdot M \cdot \frac{|\pi^*_\epsilon|}{M} + 2(\frac{d \cdot r}{d+1}) = \delta \cdot |\pi^*_\epsilon| + 2(\frac{d \cdot r}{d+1})$, where $2 \cdot (\frac{d \cdot r}{d+1})$ is the expected additive cost for connecting the endpoints of the produced solution to the query points. □

This model is shown in Figure 6, using the Monte Carlo simulation data for the value of δ, shown in Figure 4. The model is used to estimate length of an entire path, and the Monte Carlo experiments generate a sample in each of M balls and the total path length is calculated.

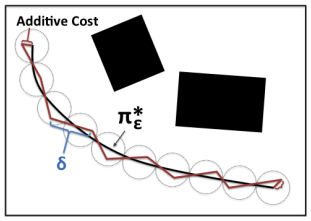

Fig. 5. Path cost for a path covering the optimum is multiplicative in δ, and incurs an additive cost.

D. A discussion of RRG and RRT*

RRG and RRT* should have similar probably near-optimal properties but their coverage process is more complex, being a frontier expansion rather than uniformly random [7]. The computation of P_{cover} in this case will also depend on the shape of the tree and of the free space it is constructed in. A possible work-around is focusing on slightly modified versions of the algorithms, where the sampling procedure can only generate points which are no more than $\frac{\epsilon}{2}$ from existing vertices in the planning structure. This should minimize the effect of the free space's shape.

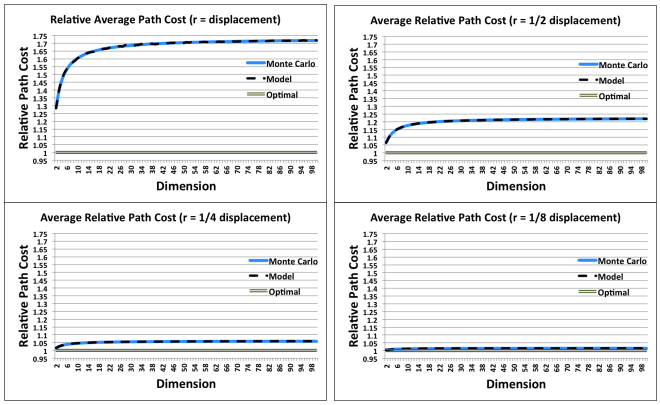

Fig. 6. Monte Carlo simulations for path distance for a problem instance where $M = 95$ and $|\pi^*_\epsilon| = 94$. Note that noise in the model is due to noise in the computed δ value. Experimental data is generated by observing total path costs, as opposed to local path costs as in Figure 4.

E. Finite-Time Properties of Roadmap Spanners

Roadmap spanner techniques (e.g., SRS [14] and IRS [13]) provide asymptotically near-optimal planning structures, which are relatively sparse. Because PNO-PRM* enforces a minimum connection radius, the resulting structure grows prohibitively large, due to its density, with expected $n^2 \cdot \frac{B(r)}{C_{\text{free}}}$ edges. Roadmap spanners are a promising direction in maintaining probably near-optimal properties, while providing the benefits of a sparse structure.

1) Worst-Path Analysis:

Theorem *3 (Worst Case Degradation for Spanners):* With probability P_{cover}, the SRS and IRS algorithms return a path of length $\ell \leq t \cdot ((\frac{\epsilon}{2} + 2r) \cdot |\pi^*_\epsilon|) + 2r$.

Proof The output of the SRS and IRS algorithms correspond to roadmap spanners of the output of PRM*. P_{cover} remains unchanged; however, it is no longer guaranteed that all consecutive sample pairs share an edge. The algorithms consider each candidate edge, and by extension all edges between consecutive sample pairs, and ensure that there exists a path through the roadmap, which is no longer than t times the length of that edge. Thus, following Theorem 1, SRS and IRS cannot produce paths more than t times longer than PRM*. □

2) Average-Path Analysis: One practically beneficial but confounding property of roadmap spanners is that paths through the spanner have much lower cost than theoretical analysis suggests. This is not well understood, making modeling path length difficult in this case.

F. Extensions

1) Stopping Criterion for PRM *methods:* Traditionally, the PRM algorithm has no well-defined stopping criterion, instead relying on ad-hoc criteria for termination. This work proposes the following natural stopping criterion for PRM by performing algebraic manipulation on Lemma 1:

Lemma *4 (A Probabilistic Stopping Criterion for* PRM*):* The PRM algorithm can be stopped after n iterations to achieve a p probability of returning a solution within the bounds drawn in Theorems 1 and 2, where

$$n = \left\lceil \frac{log(1 - p^{\frac{1}{M}})}{log(1 - \frac{|B|}{|C_{\text{free}}|})} \right\rceil$$

2) Solution Non-existence: Probabilistic completeness is not a sufficient condition to determine solution non-existence, though guarantees on solution non-existence have been studied [15], [19]; however, Lemma 1 yields a probabilistic estimate of solution non-existence.

Lemma *5 (Solution Non-Existence):* For some path, $\pi \in C_{\text{free}}$, such that $\pi(0) = q_{start}$, $\pi(1) = q_{goal}$, of length ℓ, where $d(q_{start}, q_{goal}) \leq \ell \leq \infty$, with probability P_{cover} there does not exist a solution of the given ℓ and ϵ if the algorithm has not returned a solution.

3) Extension to Non-Euclidean Spaces: A concern is whether the drawn properties hold in non-Euclidean spaces. Assume a sufficiently small ϵ, and that the space is locally homeomorphic to Euclidean space. Then, there are both local and global properties, which must be examined. Local properties are outline in lemmas 1 - 3, while the global properties are the combination of these, such as for Theorem 2.

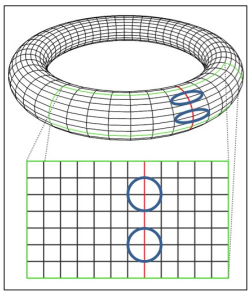

Fig. 7. Torus space, $\mathbb{T}^2$, is locally homeomorphic to $\mathbb{R}^2$. If the distance metric uses differences in angles and sampling is uniform over angular coordinates, then the properties still hold.

The critical elements of the local properties are the assumptions on the distance metric and random sampler. The local homeomorphism to Euclidean space is helpful here. For example, consider the torus, $\mathbb{T}^2$ as illustrated in Figure 7. It is straightforward to see that if the distance metric is reasoning over angle differences, then the spheres on the surface of the torus are regular in $\mathbb{R}^2$. Similarly, the sampling should also reason over angle differences, rather than attempting to sample uniformly on the surface of the torus embedded in $\mathbb{R}^3$. Any space which is locally homeomorphic to Euclidean space shares similar properties, but this is not a general result for all spaces.

When examining the global properties, it is easy to see that despite the unifications in the space, optimal paths will still be covered by a set of disjoint balls. The local neighborhood of those balls still being homeomorphic to Euclidean ensures connections will be made and eventually a covering path can be constructed.

V. Experimental Validation

Experiments were run in the Open Motion Planning Library (OMPL) for point systems [1]. Experiments were performed in six environments: no obstacles ($\mathbb{R}^2$, $\mathbb{T}^2$, $SE(2)$, and $SE(3)$), 2D corridor, and 3D cube.

A. Probability of Path Coverage

Verifying the accuracy of P_{cover} is critical to the entire formulation. Table I shows the automated stopping criterion thresholds for the environments for confidence intervals 0.90, 0.95, and 0.99 (top), and the ratio of experiments which completed within this limit (bottom).

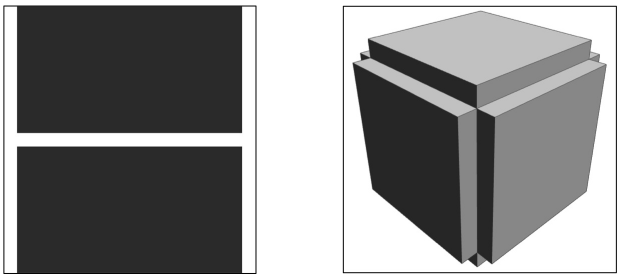

Fig. 8. Environments with obstacles are the Corridor (left) and Cube (right).

Environment	$n_{0.90}$	$n_{0.95}$	$n_{0.99}$	M	ϵ
Obstacle-Free	2368	2733	3561	11	1.0
Corridor	662	764	995	11	1.0
Cube	7390	8529	11115	11	1.0
Torus	5845	6746	8792	11	0.4
SE(2)	89337	103110	134368	11	1.0
SE(3)	2891429	3402015	4562985	6	2.0

Environment	$p = 0.90$	$p = 0.95$	$p = 0.99$
Obstacle-Free	0.917	0.954	0.985
Corridor	0.898	0.947	0.991
Cube	0.900	0.954	0.991
Torus	0.895	0.951	0.991
SE(2)	0.916	0.965	0.992
SE(3)	0.891	0.952	0.986

TABLE I

EXPERIMENTAL TESTS FOR COVERING π^*_ϵ.

These results show good correlation with the analysis. The ratio for $\frac{|B(r)|}{|C_{\text{free}}|}$ was approximated through Monte Carlo experiments. The approximation has error, which does affect P_{cover}. This ratio could be computed during the construction of the roadmap; however, this will also increase the error, which could adversely affect P_{cover}, as the Monte Carlo experiments tested billions of samples for high accuracy.

B. Path Cost

The analysis also provides bounds on path length, obtained with probability P_{cover}. Figure 9 shows lengths of average paths extracted from the planning structure. Note that the data points were taken as soon as coverage of π^*_ϵ was obtained. Tests in $SE(2)$ and $SE(3)$ deviate from the drawn bound, likely due to the δ value having been computed for spheres in Euclidean space.

VI. CONCLUSION

The finite-time properties of sampling-based planners are examined, and bounds are drawn for expected path cost and probability of attaining a solution. The analysis allows a natural stopping criterion for sampling-based planners to produce probably near-optimal solutions after finite computation. The models drawn are verified against Monte Carlo experiments and traditional planning methods.

This work opens many avenues for further investigation. Critically important is developing a model for "returned" paths, and extending the work for tree-based planners. Due to the limitations of PNO-PRM*, it is pertinent to draw better bounds for roadmap spanner techniques. It would also be better to have more reliable and robust models for worst-case bounds and the required δ value for average-case analysis even for non-Euclidean spaces.

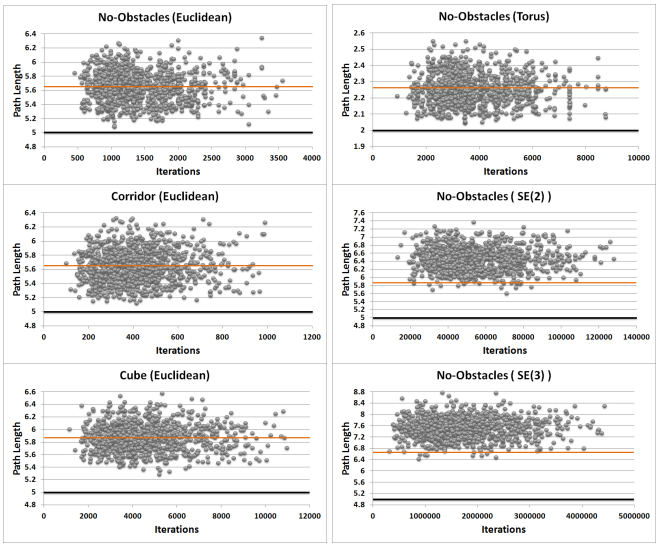

Fig. 9. Path lengths at various iterations for average paths extracted from the planning structure. The black line shows the optimal path length, and the orange is the expected length.

REFERENCES

[1] I. A. Şucan, M. Moll, and L. E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012.

[2] A. Dobson and K. E. Bekris. Improving Sparse Roadmap Spanners. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2013.

[3] A. Dobson, T. D. Krontiris, and K. E. Bekris. Sparse Roadmap Spanners. In *WAFR*, Cambridge, MA, June 2012.

[4] D. Hsu, L. Kavraki, J.-C. Latombe, R. Motwani, and S. Sorkin. On Finding Narrow Passages with Probabilistic Roadmap Planners. In *WAFR*, Houston, TX, 1998.

[5] L. Jaillet and T. Simeon. Path Deformation Roadmaps. In *WAFR*, New York City, NY, July 2006.

[6] S. Karaman and E. Frazzoli. Incremental Sampling-based Algorithms for Optimal Motion Planning. In *RSS*, Zaragoza, Spain, 2010.

[7] S. Karaman and E. Frazzoli. Sampling-based Algorithms for Optimal Motion Planning. *IJRR*, 30(7):846–894, June 2011.

[8] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe. Analysis of Probabilistic Roadmaps for Path Planning. *IEEE TRA*, 14(1):166–171, 1998.

[9] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars. Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces. *IEEE TRA*, 12(4):566–580, 1996.

[10] M. Kearns, Y. Mansour, and A. Y. Ng. A sparse sampling algorithm for near-optimal planning in large markov decision processes. *MACHINE LEARNING*, 49:193–208, 2002.

[11] A. M. Ladd and L. E. Kavraki. Measure Theoretic Analysis of Probabilistic Path Planning. *IEEE TRA*, 20(2):229–242, April 2004.

[12] S. M. LaValle and J. J. Kuffner. Randomized Kinodynamic Planning. *IJRR*, 20:378–400, May 2001.

[13] J. D. Marble and K. E. Bekris. Asymptotically Near-Optimal is Good Enough for Motion Planning. In *ISRR*, Flagstaff, AZ, August 2011.

[14] J. D. Marble and K. E. Bekris. Computing Spanners of Asympotically Optimal Probabilistic Roadmaps. In *IEEE/RSJ IROS*, San Francisco, CA, September 2011.

[15] Z. McCarthy, T. Bretl, and S. Hutchinson. Proving path non-existence using sampling and alpha shapes. In *ICRA*, May 2012.

[16] B. Raveh, A. Enosh, and D. Halperin. A Little More, a Lot Better: Improving Path Quality by a Path-Merging Algorithm. *IEEE TRO*, 27(2):365–370, 2011.

[17] L. A. Santalo. *Integral Geometry and Geometric Probability*, volume 1 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley Publishing Company, Reading, Massachusets, 1976.

[18] L. G. Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[19] G. Varadhan and D. Manocha. Star-shaped Roadmaps: A Deterministic Sampling Approach for Complete Motion Planning. *RSS*, 2005.