Geometric Probability Results for Bounding Path Quality in Sampling-based Roadmaps after Finite Computation

Andrew Dobson[1] George V. Moustakides[2] Kostas E. Bekris[1]

***Abstract*— Sampling-based algorithms provide efficient solutions to high-dimensional, geometrically complex motion planning problems. For these methods asymptotic results are known in terms of completeness and optimality. Previous work by the authors argued that such methods also provide probabilistic near-optimality after finite computation time using indications from Monte Carlo experiments. This work formalizes these guarantees and provides a bound on the probability of finding a near-optimal solution with PRM* after a finite number of iterations. This bound is proven for general-dimension Euclidean spaces and evaluated through simulation. These results are leveraged to create automated stopping criteria for PRM* and sparser near-optimal roadmaps, which have reduced running time and storage requirements.**

I. Introduction and Related Work

While motion planning is PSPACE-hard [1], [2], sampling-based approaches can frequently provide efficient solutions in practice, even for high-dimensional, geometrically complex problems [3], [4]. Two families of methods have emerged. Approaches, such as the Probabilistic Roadmap Method (PRM), preprocess a configuration space ($\mathcal{C}$-space) to create a structure useful for answering multiple queries [5]. Tree-based planners, such as the Rapidly-exploring Random Tree (RRT), are suited for single-query planning and kinodynamic systems [6], [7]. Many variants focus on quickly finding solutions and dealing with narrow passages [8]. Some approaches are tailored to returning high-clearance paths [9], or low cost solutions [10].

Formal analysis showed that sampling-based planners provide probabilistic completeness, i.e., if the instance has a solution, the methods solve the problem with probability approaching one [11], [12]. In the general case, solution non-existence cannot be detected, although there are efforts in this direction [13]. Recent progress identified the conditions under which sampling-based approaches, such as PRM* and RRT*, achieve asymptotic optimality [14]. This property ensures that the computed paths asymptotically converge to the optimal ones.

Asymptotic results do not characterize solution quality after a finite time of computation. Experimental trials show that asymptotically optimal planners provide good quality paths in practice, even if optimality constraints are relaxed [15], [16], [17], [18], [19]. To bridge the gap between theory and practice, bounds on path length after finite computation are needed.

This work formalizes the *Probabilistic Near-Optimality* (PNO) of efficient sampling-based roadmaps under limited assumptions. A previous contribution by the authors showed PNO properties for a PRM variant, which asymptotically converges to a dense planning structure, relying on Monte Carlo simulations to draw bounds on path length and estimate parameters of the bound [20]. Related work has initiated a similar study for the FMT* algorithm [21]; however, the bound provided in the FMT* work does not result in an automated stopping criterion, as there are several free parameters. The current paper provides the following contributions:

- A novel approximation to an unsolved problem in geometric probability (to the best of the authors' knowledge) to draw closed-form path length bounds.
- An analysis for PRM* as opposed to the previously considered PNO-PRM*, which results in a denser roadmap [20].

The analysis reasons over a construction of hyperballs, which cover optimal paths in the $\mathcal{C}$-space. Experimental validation using simulations accompanies this analysis.

PNO properties can impact the practical use of motion planning methods. For example, in robot task planning, motion planners are queried frequently to determine a solution to perform a higher-level task [22]. PNO properties expose information, such as degradation from optimal along a given homotopic class, and bounds on solution non-existence, which can be useful for pruning the search along different homotopic classes. This work also results in an automated stopping criterion, which is informed in terms of solution quality.

Work by the authors has been supported by NSF CCF:13307893, NSF IIS:1451737 and a DHS scholarship to Andrew Dobson. Any conclusions expressed here are of the authors and do not reflect the views of the sponsors.

[1]Computer Science, Rutgers University, NJ, USA, {ajd223,kostas.bekris}@cs.rutgers.edu.

[2]Electrical and Computer Engineering, University of Patras, Greece, moustaki@upatras.gr.

II. Problem Setup

This work examines kinematic planning in the configuration space $\mathcal{C}$, the set of robot configurations $q \in \mathcal{C}$. $\mathcal{C}$ is partitioned into the collision-free ($\mathcal{C}_{\text{free}}$) and colliding ($\mathcal{C}_{\text{obs}}$) subsets. This work reasons over $\mathcal{C}$ as a metric space, using the Euclidean L_2-norm as a distance metric. For robustly feasible motion planning problems, there exists a set of δ-robust paths, i.e., paths with a minimum clearance with $\mathcal{C}_{\text{obs}}$ of at least δ, which answer a query (q_{start}, q_{goal}). Let the path of minimum length from that set be denoted as $\pi^*_{\delta_n}$ with length $L^*_{\delta_n}$. Then, this work examines the following problem:

Defn. 1 (Robustly Feasible Motion Planning): For a Robustly Feasible Motion Planning (RFMP) Problem $(\mathcal{C}, q_{start}, q_{goal}, \delta_{n_0})$ such that an optimal δ_{n_0}-robust path $\pi^*_{\delta_n}$ exists, where $\pi^*_{\delta_n}(0) = q_{start}$ and $\pi^*_{\delta_n}(1) = q_{goal}$, find a solution path $\pi : [0,1] \rightarrow \mathcal{C}_{\text{free}}$, so that $\pi(0) = q_{start}$ and $\pi(1) = q_{goal}$.

Algorithm 1: PRM*(n)

```
V ← ∅; E ← ∅;
for i = 1 ... n do
    v ← SAMPLE_FREE();
    V ← V ∪ v;
    r_n ← CONNECT_RADIUS(i);
    U ← NEAR(V, v, r_m);
    for u ∈ U do
        if L(v, u) ∈ C_free then
            E ← E ∪ {L(v, u)};
return G = (V, E);
```

This paper studies a variant of PRM*, which is an asymptotically optimal method that results in a sparser roadmap than PNO-PRM* considered in the prior work by the authors [20]. The high-level operations of the PRM* framework are outlined in Algorithm 1. The difference lies in the radius used in the subroutine CONNECT_RADIUS. The considered variant uses twice the connection radius of the original PRM* ($r_n = 2 \cdot r_{\text{PRM*}}$) to achieve the following property:

Prop. 1 (Probabilistic Near-Optimality for RFMP): An algorithm ALG is probabilistically near-optimal for an RFMP problem $(\mathcal{C}, q_{start}, q_{goal}, \delta_{n_0})$, if for a finite iteration $n \geq n_0$ of ALG and an error threshold ϵ, there is a probability $\mathbb{P}_{success}$ that ALG returns a solution path π_n of length L_n, such that:

$$\mathbb{P}(|L_n - L^*_{\delta_n}| > \epsilon \cdot L^*_{\delta_n}) < 1 - \mathbb{P}_{success},$$

where $L^*_{\delta_n}$ is the length of the optimum δ_n-robust path $\pi^*_{\delta_n}$ for a clearance value $\delta_n < \delta_{n_0}$, and n_0 is a minimum iteration that this guarantee can be provided.

Computing a bound n_0 requires a pessimistic estimate of the length of the optimal path $L^*_{\delta_n}$. The analysis proceeds with the following high-level steps: first, some basic definitions and constructions are given, and then the probability of the method constructing a specific class of path is provided. Next, the expected value and variance of the length of such a path is approximated, and finally, a probabilistically near-optimal bound is derived using the Chebyshev inequality.

For the reader's convenience, a glossary of terms is available here:

Term	Definition
δ_n	Clearance of optimum $\pi^*_{\delta_n}$ at iteration n.
ϵ	Multiplicative error on path length.
L_n	Rand. Var. for length of returned path.
$L^*_{\delta_n}$	Rand. Var. for length of optimal path.
$\mathbb{P}_{success}$	Probability of generating a low-error path.
n_0	Iteration where PNO requirements are met.
$\mathcal{B}_\rho(x)$	Hyperball of radius ρ centered at x.
M_n	Number of hyperballs at iteration n.
ρ_n	Radius of the covering hyperballs.
$\mathbb{P}(\Omega^\rho_n)$	Probability of generating a covering path.

III. Analysis Preliminaries

The analysis considers a theoretical construction of hyperballs as illustrated in Figure 1. Consider $M_n + 1 = \lceil \frac{L^*_{\delta_n}}{\delta_n} \rceil + 1$ balls $\mathcal{B}_n$ centered along $\pi^*_{\delta_n}$, i.e.

$$\mathcal{B}_n = \{\mathcal{B}_{\rho_n}(\pi^*_{\delta_n}(\tau_0)), \ldots, \mathcal{B}_{\rho_n}(\pi^*_{\delta_n}(\tau_{M_n}))\},$$

having radius $\rho_n \leq \frac{1}{2}\delta_n$, where $\delta_n = \frac{1}{2}r_n$, and where r_n is the connection radius used by this PRM* variant. These disjoint hyperballs are centered δ_n distance apart. Any pair of points lying in consecutive hyperballs is tested for connection by PRM*. This construction is similar to the analysis of PRM*, when $\theta_1 = 2$ (Definition 51) [14].

Consider for each hyperball, the set of samples S generated by PRM*, which lie in this hyperball. If $|S| \geq 1$ for all of the hyperballs, then there must exist at least one path through these hyperballs in the PRM* graph. The analysis considers the existence and length L_n of this path in order to provide PNO guarantees.

To ensure this path is constructed, the appropriate connection radius r_n must be used, which in the presentation of the PRM*algorithm it is a function of a parameter γ_{PNO}. It is also possible to consider the use of a k_{PNO}-PRM* algorithm, where at each iteration the new sample is connected to the closest k_{PNO} neighbors. Parameter k_{PNO} matches the expected number of samples in the radius $r_n(\gamma_{\text{PNO}})$. The values γ_{PNO} and k_{PNO} are given as Lemmas below for brevity purposes. Full derivations are provided in an extended version of this work [23].

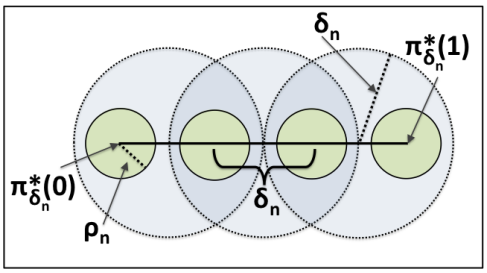

Fig. 1. Hyperballs over an optimal path with radius ρ_n and separation δ_n. Consecutive balls lie entirely within a hyperball of radius δ_n centered at τ_t: $\mathcal{B}_{\delta_n}(\pi^*_{\delta_n}(\tau_t))$.

Lemma 1 (Connectivity constant $\gamma_{\texttt{PNO}}$): To ensure PNO properties of PRM*, it suffices to use $\gamma_{\texttt{PNO}}$ parameter:

$$\gamma_{\texttt{PNO}} > 4\left(\left(1+\frac{1}{d}\right)\left(\frac{|\mathcal{C}_{\text{free}}|}{V_d}\right)\right)^{\frac{1}{d}} = 2\cdot\gamma_{\texttt{PRM}^*}$$

Lemma 2 (Connectivity constant $k_{\texttt{PNO}}$): To ensure PNO properties of k-PRM*, it suffices to use:

$$k_{\texttt{PNO}} = 2^d$$

Here, $|\mathcal{C}_{\text{free}}|$ indicates the volume of $\mathcal{C}_{\text{free}}$, and $V_d = |\mathcal{B}_1(\cdot)|$ is the volume of a d-dimensional unit hyperball.

At a high level, the reason the radius doubles relatively to the analysis of the original PRM* algorithm follows from the difference in hyperball constructions. While using balls of the same radius ρ_n, PRM*'s analysis considered the case where the separation of the balls is 0. The current work has a separation of $\delta_n \geq 2\cdot\rho_n$. Similarly, the previous analysis of PNO-PRM* differs in that it used a fixed number of hyperballs [20]. The current work parameterizes the separation and radius of the hyperballs as functions over the number of drawn samples n.

IV. Probability of Path Coverage

Foundational work on roadmap-based approaches sought to characterize how quickly these methods return valid solutions [11]. An exponential bound on this probability was derived as:

$$\mathbb{P}(\Omega^\rho_n) \geq 1 - \frac{2L^*_{\delta_n}}{\delta_n} e^{-\frac{\pi\cdot\delta_n^2\cdot n}{4|\mathcal{C}_{\text{free}}|}} \tag{1}$$

Here, $\mathbb{P}(\Omega^\rho_n)$ represents the probability that a path has been generated in a set of hyperballs centered around a clearance-robust optimal path. A tighter bound was derived in prior work for PNO-PRM* [20], while this work presents a similarly tight bound for PRM*. As an example, for a two-dimensional configuration space with $|\mathcal{C}_{\text{free}}| = 100$, $L^*_{\delta_n} = 10$, $\delta_n = 1$, and for $\mathbb{P}(\Omega^\rho_n) = 0.99$, the derived bound accurately reports 97 samples are required as opposed to nearly 980 with Equation 1. This bound follows from prior work in the literature [11], [14], [20]. The derivation of $\gamma_{\texttt{PNO}}$ also gives the value for the hyperballs' radii, $\rho_n = \left(\left(1+\frac{1}{d}\right)\left(\frac{|\mathcal{C}_{\text{free}}|}{V_d}\right)\left(\frac{\ln n}{n}\right)\right)^{\frac{1}{d}}$, as well as the number of covering hyperballs, $M_n + 1 = \lceil\frac{L_n}{2\rho_n}\rceil + 1$. Substituting these values into the existing coverage probability results from related work yields:

$$\mathbb{P}(\Omega^\rho_n) \approx \left(1-\left(1-a\right)^n\right)^{\frac{1}{2}L^*_{\delta_n}\left(b\right)^{-\frac{1}{d}}+1}$$

where $a = \frac{V_d\left(\sqrt[d]{\left(1+\frac{1}{d}\right)\left(\frac{|\mathcal{C}_{\text{free}}|\ln n}{nV_d}\right)}\right)^d}{|\mathcal{C}_{\text{free}}|}$, $b = \left(1+\frac{1}{d}\right)\left(\frac{|\mathcal{C}_{\text{free}}|\ln n}{nV_d}\right)$. The value of a arises from the relative size of the hyperballs in the free space, while b is related to M_n. Simplifying this expression yields the following Lemma:

Lemma 3 (Probability of Path Coverage): Let Ω^ρ_n be the event that for one execution of PRM* there exists at least one sample in each of the $M_n + 1$ hyperballs of radius ρ_n over the clearance robust optimal path, $\pi^*_{\delta_n}$, for a specific value of $n > n_0$ and ρ_n. Then,

$$\mathbb{P}(\Omega^\rho_n) \approx \left(1-\left(1-a\right)^n\right)^{\frac{1}{2}L^*_{\delta_n}\left(b\right)^{-\frac{1}{d}}+1} \tag{2}$$

where $a = \left(1+\frac{1}{d}\right)\left(\frac{\log n}{n}\right)$ and $b = \frac{|\mathcal{C}_{\text{free}}|}{V_d}a$.

V. Bounding Path Quality

The high-level approach to determining the path quality bound is done in four steps. It first expresses the bound in terms of the mean and variance of L_n. Then, the expected value and variance of L_n are derived. Finally, the results are combined.

A. Deriving a bound in terms of mean and variance

Let Ω^C be the event that there does not exist a sample in each of the hyperballs covering a path, i.e., $\mathbb{P}(\Omega^C) = 1 - \mathbb{P}(\Omega^\rho_n)$. Then, the value for $\mathbb{P}(|L_n - L^*_{\delta_n}| > \epsilon\cdot L^*_{\delta_n})$ can be expressed as:

$$\begin{aligned}\mathbb{P}\big(|L_n - L^*_{\delta_n}| > \epsilon\cdot L^*_{\delta_n} \mid \Omega^\rho_n\big)\mathbb{P}(\Omega^\rho_n)\\ + \mathbb{P}\big(|L_n - L^*_{\delta_n}| > \epsilon\cdot L^*_{\delta_n} \mid \Omega^C\big)\mathbb{P}(\Omega^C)\end{aligned}$$

It is assumed that $\mathbb{P}\big(|L_n - L^*_{\delta_n}| > \epsilon\cdot L^*_{\delta_n} \mid \Omega^C\big)$ is close to 1 and is thus upper bounded by 1. Let y be a random variable identically distributed with L_n, but having zero mean, i.e., $y = L_n - \mathbb{E}[L_n]$. Then,

$$\begin{aligned}&\mathbb{P}(|L_n - L^*_{\delta_n}| > \epsilon\cdot L^*_{\delta_n}) =\\ &\quad\mathbb{P}\big(\mathbb{E}[L_n] + y - L^*_{\delta_n} > \epsilon\cdot L^*_{\delta_n} \mid \Omega^\rho_n\big)\\ &\quad+\mathbb{P}\big(\mathbb{E}[L_n] + y - L^*_{\delta_n} < -\epsilon\cdot L^*_{\delta_n} \mid \Omega^\rho_n\big),\end{aligned}$$

Then, exploiting symmetry and some algebra:

$$\begin{aligned}&\mathbb{P}\big(|L_n - L^*_{\delta_n}| > \epsilon\cdot L^*_{\delta_n} \mid \Omega^\rho_n\big) =\\ &\qquad 2\mathbb{P}\big(y > (\epsilon+1)\cdot L^*_{\delta_n} - \mathbb{E}[L_n] \mid \Omega^\rho_n\big)\end{aligned}$$

Bound the right-hand side with Chebyshev's Inequality:

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq a) \leq \frac{Var(X)}{a^2}$$

In order to employ this inequality, both $\mathbb{E}[L_n]$ and $Var(L_n)$ are needed.

B. Approximation of $\mathbb{E}[L_n]$ in $\mathbb{R}^d$

For the sake of brevity, only the high-level approach for deriving $\mathbb{E}[L_n]$ is given.

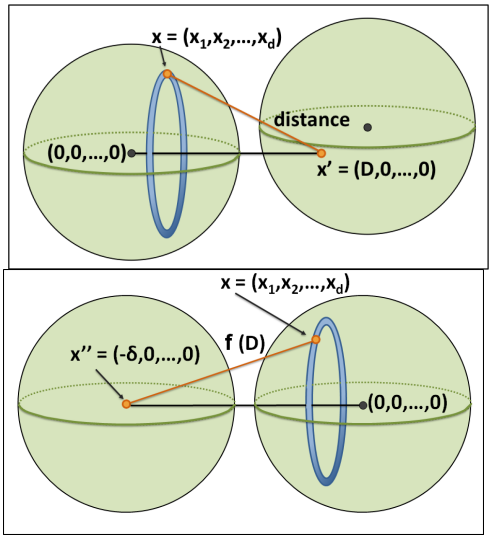

Fig. 2. The mean calculation in $\mathbb{R}^3$. (top) The first set of integrals is performed over the left hyperball, averaging the distance between points $(x_1, x_2, \ldots, x_d)$ and $(D, 0, \ldots, 0)$. (bottom) A second set of integrals is performed over the second hyperball using the above result, yielding the expected value.

Let, $\mathbb{E}[L_n] = \sum_{m=1}^{M_n} \mathbb{E}[L_m]$, where L_m is the length of a single segment between hyperballs. Then, because all L_m are I.I.D., $\mathbb{E}[L_n] = M_n \mathbb{E}[L_1]$. Then, $\mathbb{E}[L_1]$ is approximated via two integrations over the endpoints of the segment, as outlined in Figure 2. The first integration uses D, the distance between the center of the first hyperball and a point in the second, while the second integration integrates over the function, $f(D)$. The first integration attains the following form:

$$A = \frac{1}{V_d \rho_n^d} \int \cdots \int_{x_1^2 + \ldots + x_d^2 \leq \rho_n^2} \sqrt{(x_1 - D)^2 + x_2^2 + \ldots + x_d^2}\, dx_1 \ldots dx_d,$$

A second-order Taylor Approximation is employed, assuming that ρ_n much smaller than δ_n. This integral can be simplified using the following Lemmas, stated without proof:

Lemma 4 (Recurrence relation of $\int_0^\pi (\sin\theta)^d d\theta$): For $S_d = \int_0^\pi (\sin\theta)^d d\theta$, the following recurrence relation holds:

$$\int_0^\pi (\sin\theta)^d d\theta = S_d = \frac{d-1}{d} S_{d-2}$$

Lemma 5 (Value of $\int_0^\pi (\sin\theta)^d d\theta$): In terms of the hyperball volume constant, V_d,

$$\int_0^\pi (\sin\theta)^{d-2} d\theta = S_{d-2} = \frac{dV_d}{(d-1)V_{d-1}}$$

These Lemmas lead to the following intermediate form:

$$A = D + \frac{(d-1)\rho_n^2}{2(d+2)D}. \tag{3}$$

Integrate over A to attain $\mathbb{E}[L_n]$. The steps taken are similar to the steps taken above, as the form of the integral nearly the same. After simplifying, the following Lemma arises:

Lemma 6 (Expected value of L_n): The path built over the set of $M_n + 1$ hyperballs having radius ρ_n has expected length:

$$\mathbb{E}[L_n] \approx M_n \Big(\delta_n + \frac{(d-1)\rho_n^2}{(d+2)\delta_n}\Big)$$

To verify the approximation, Monte Carlo experiments were employed. The relative error of the approximation to the simulated values are shown in Figure 3:

Euclidean	$\lambda_n = 0.5$	$\lambda_n = 0.125$
dimension	% error	% error
2	0.1730%	0.0050%
3	0.0473%	0.0205%
10	0.9413%	0.0128%
100	1.9147%	0.0129%

Fig. 3. Simulation comparison for $\mathbb{E}[L_n]$, using 120,000 data points for each entry, for differing $\lambda_n = \frac{\rho_n}{\delta_n}$, where the error is $100 \cdot \frac{|\mathbb{E}[L_n] - L_n|}{L_n}$.

C. Computation of the Variance of L_n in $\mathbb{R}^d$

To compute the $Var(L_n)$, leverage the definition of the variance of a random variable, i.e. $Var(X) = \mathbb{E}[X^2] - (\mathbb{E}[X])^2$:

$$Var\Big(\sum_{m=1}^{M_n} L_m\Big) = \mathbb{E}\Big[\sum_{m=1}^{M_n} L_m^2\Big] - \Big(\mathbb{E}\Big[\sum_{m=1}^{M_n} L_m\Big]\Big)^2$$
$$= \sum_{m=1}^{M_n} \sum_{k=1}^{M_n} \mathbb{E}[L_m L_k] - \Big(\mathbb{E}\Big[\sum_{m=1}^{M_n} L_m\Big]\Big)^2$$

The second term can be simplified due to the linearity of expectation:

$$Var\big(\mathbb{E}[L_m L_k] - M_n^2 \big(\mathbb{E}[L_m]\big)^2$$

Here, only the variance terms of each segment and covariance between adjacent segments contribut to the sum, allowing a simplification to

$$Var\big(\sum_{m=1}^{M_n} L_m\big) = M_n\mathbb{E}[L_1^2] + (2M_n-2)\mathbb{E}[L_1L_2] + (2-3M_n)\big(\mathbb{E}[L_1]\big)^2$$

$\mathbb{E}[L_1^2]$ and $\mathbb{E}[L_1L_2]$ are unknown. For the sake of brevity, derivations for these values are omitted, and their values are given.

Lemma 7 (Expected value of L_1^2): For two consecutive hyperballs, the expected squared distance between random points in those balls is

$$\mathbb{E}[L_1^2] = \delta_n^2 + \frac{2d}{d+2}\rho_n^2$$

Lemma 8 (Expected value of L_1L_2 in $\mathbb{R}^d$): For three consecutive hyperballs, the expected value of the product of the lengths of the segments connecting random samples inside those balls is

$$\mathbb{E}[L_1L_2] \approx \delta_n^2 + \Big(\frac{2d-3}{d+2}\Big)\rho_n^2$$

Subsitute the values from Lemmas 6, 7, and 8 into the above form to get:

$$Var\big(\sum_{m=1}^{M_n} L_m\big) \approx M_n\Big(\delta_n^2 + \frac{2d}{d+2}\rho_n^2\Big) + (2M_n-2)\Big(\delta_n^2 + \frac{2d-3}{d+2}\rho_n^2\Big) + (2-3M_n)\Big(\delta_n + \frac{(d-1)\rho_n^2}{(d+2)\delta_n}\Big)^2$$

After simplification, the following Lemma can be reached:

Lemma 9 (Variance of L_n): L_n has variance:

$$Var(L_n) \approx \frac{2\rho_n^2}{d+2}$$

Simulations verify that the drawn approximation characterizes the variance properly, as shown in Figure 4:

Euclidean	$\lambda_n = 0.5$	$\lambda_n = 0.125$
dimension	% error	% error
2	6.0245%	0.5739%
3	9.7691%	1.0655%
10	19.0989%	2.1429%
100	23.7279%	2.8191%

Fig. 4. Simulation comparison for $Var(L_n)$, using 120,000 data points for each entry, where the error is $100 \cdot \frac{|Var(L_n) - Var^{MC}|}{Var^{MC}}$.

D. Finalizing the PNO guarantee of PRM*

Using the computed mean and variance, the final bound is derived. Substituting into the form reached in Section V-A yields:

$$2\mathbb{P}\big(y > M_n\delta_n\big(\delta_n - \frac{(d-1)\rho_n^2}{(d+2)\delta_n^2}\big) \mid \Omega_n^\rho\big)$$

Applying Chebyshev's Inequality results in the following theorem:

Thm. 1 (Probabilistic Near-Optimality of PRM):* For finite iterations n, PRM* is probabilistically near-optimal, returning a path of length L_n such that

$$\mathbb{P}\big(|L_n - L^*_{\delta_n}| \geq \epsilon \cdot L^*_{\delta_n}\big) \leq 1 + \mathbb{P}\big(\Omega_n^\rho\big)(\chi - 1),$$
$$\text{where } \chi = \frac{\frac{4\lambda_n^2\delta_n^2}{d+2}}{{L^*_{\delta_n}}^2\big(\epsilon - \frac{(d-1)}{(d+2)}\lambda_n^2\big)^2} \quad (4)$$

VI. Using PNO properties in practice

This section highlights various ways PNO properties can be leveraged in practice.

A. Extending PNO to roadmap spanners.

As a matter of practicallity, PNO properties extend naturally to certain practical methods which reduce memory requirements [17]. In particular, two such methods, SRS and IRS, provide a spanner over the output of PRM*, which leads to the following Corollary:

Corr. 1 (PNO of SRS and IRS): For finite iterations n and input stretch t, SRS and IRS are PNO, probabilistically containing a path of length L_{span} such that

$$\mathbb{P}\big(|L_{span} - t \cdot L^*_{\delta_n}| \geq \epsilon \cdot t \cdot L^*_{\delta_n}\big) \leq 1 + \mathbb{P}\big(\Omega_n^\rho\big)(\chi - 1)$$

B. Online Prediction of $L^*_{\delta_n}$

The length of the optimal path in the same homotopic class as PRM*'s current solution can be approximated. Estimate $L^*_{\delta_n}$ by considering the number of hyperballs M_n+1 to estimate ϵ. Then, use the current returned path length from the algorithm, L_n, and set $L^*_{\delta_n} = \frac{L_n}{(\epsilon+1)}$. From the analyis, it can be shown that:

$$\mathbb{P}\big(L_n - L^*_{\delta_n} \geq \epsilon \cdot L^*_{\delta_n}\big) = \mathbb{P}\big(L^*_{\delta_n} \leq \frac{L_n}{\epsilon+1}\big) \leq \mathbb{P}\big(|L_n - L^*_{\delta_n}| \geq \epsilon \cdot L^*_{\delta_n}\big) \leq 1 + \mathbb{P}\big(\Omega_n^\rho\big)(\chi - 1)$$

Consider, however, that this result is only valid given that $\pi^*_{\delta_n}$ exists for the current value of δ_n. Therefore, it is critical that the algorithm executes at least until $\delta_n \leq \delta_{n_0}$, i.e. when $n \geq n_0$. Then all that remains is to solve the bound in terms of ϵ. From above:

$$\mathbb{P}\big(L^*_{\delta_n} \leq \frac{L_n}{\epsilon+1}\big) = 1 - \mathbb{P}_{success} \leq 1 + \mathbb{P}\big(\Omega_n^\rho\big)(\chi - 1)$$

Performing some algebraic manipulation yields:

$$\chi \geq 1 - \frac{\mathbb{P}_{success}}{\mathbb{P}(\Omega_n^\rho)} \quad (5)$$

Then, substituting χ and simplifying leads to the following Lemma:

Lemma 10 (Multiplicative bound ϵ_n): After $n > n_0$ iterations of PRM*, with probability $\mathbb{P}_{success}$, if $\pi^*_{\delta_n}$

exists, then PRM* contains a path ϵ_n-bounded by $L^*_{\delta_n}$ where:

$$\epsilon_n \leq \frac{2\lambda_n \delta_n}{L^*_{\delta_n}} \sqrt{\frac{1}{(d+2)(1-\frac{\mathbb{P}_{success}}{\mathbb{P}(\Omega^\rho_n)})}} + \frac{(d-1)}{(d+2)}\lambda_n^2 \quad (6)$$

Finally, estimate $L^*_{\delta_n}$ as $L^*_{\delta_n} \approx \frac{L_n}{(\epsilon_n+1)}$.

C. Deriving probabilistic stopping criteria

This section derives n_0 for a desired confidence probability $\mathbb{P}_{DES}$ within an error bound ϵ_{DES}. Let $\lambda_n = \frac{1}{2}$ and then solving Equation 6 for $\mathbb{P}(\Omega^\rho_n)$ yields:

$$(d+2)\big(1-\frac{\mathbb{P}_{success}}{\mathbb{P}(\Omega^\rho_n)}\big) \geq \frac{1}{M_n^2\big(\epsilon_n - \frac{1}{4}\frac{(d-1)}{(d+2)}\big)^2}$$

Solving for $\mathbb{P}(\Omega^\rho_n)$, the right hand side will be denoted as ψ:

$$\mathbb{P}(\Omega^\rho_n) \geq \frac{1}{\frac{1}{\mathbb{P}_{DES}} \cdot \Big(1 - \frac{1}{M^2_{n_0}\cdot(d+2)\big(\epsilon_{DES} - \frac{(d-1)}{4(d+2)}\big)^2}\Big)} = \psi$$

Then, substituting the form of Equation 2 using ρ_0, M_{n_0}, and n_0, and solving for n_0 yields the following Lemma:

*Lemma 11 (*PNO *iteration limit for* PRM*):* For given ϵ_{DES} and $\mathbb{P}_{DES}$, the graph of PRM* probabilistically contains a path π_0 of length L_0 with $\mathbb{P}\big(|L_0 - L^*_{\delta_0}| \geq \epsilon_{DES} \cdot L^*_{\delta_0}\big) \leq 1 - \mathbb{P}_{DES}$ after n_0 iterations, where

$$n_0 \leq \left\lceil \frac{\log\left(1 - \sqrt[M_{n_0}]{\psi}\right)}{\log\left(1 - \frac{|\mathcal{B}_{\rho_0}|}{|\mathcal{C}_{\text{free}}|}\right)} \right\rceil, \text{where,}$$

$$\psi = \frac{1}{\frac{1}{\mathbb{P}_{DES}} \cdot \Big(1 - \frac{1}{M^2_{n_0}\cdot(d+2)\big(\epsilon_{DES} - \frac{(d-1)}{4(d+2)}\big)^2}\Big)} \quad (7)$$

When $n \geq n_0$, if no path has been returned within the error bound, then with probability at least $\mathbb{P}(\Omega^\rho_n)$, no such optimal path exists within the supposed homotopic class of $\pi^*_{\delta_n}$. It is left to future work to determine the best approach for reasoning over multiple homotopic classes.

($\mathbb{R}^3$) n	$\epsilon = 0.1$	$\epsilon = 0.2$	$\epsilon = 0.5$	$\epsilon = 1.0$
100	0.0028	0.1424	0.7601	0.9362
1,000	0.0666	0.5815	0.9330	0.9835
100,000	0.7675	0.9635	0.9957	0.9990
10,000,000	0.9785	0.9977	0.9998	0.9999
($\mathbb{R}^6$) n	$\epsilon = 0.1$	$\epsilon = 0.2$	$\epsilon = 0.5$	$\epsilon = 1.0$
500	0.0000	0.0033	0.6335	0.8235
5,000	0.0000	0.0845	0.8584	0.9712
500,000	0.0017	0.5884	0.9669	0.9934
50,000,000	0.5914	0.8856	0.9924	0.9985

Fig. 5. Probability of returning near-optimal paths for 3D and 6D collision-free problem instances. Here, $|\mathcal{C}_{\text{free}}|$ = 1,000 ($\mathbb{R}^3$), and 216,000 ($\mathbb{R}^6$).

The tables in Figure 5 provide an indication of the required iterations for problems of varying dimension.

VII. Simulations

Simulations were performed in four environments on the PRACSYS simulation software [24]. Environments with obstacles are shown in Figure 6, which were created to be highly regular, to simplify the computation of $|\mathcal{C}_{\text{free}}|$ and $L^*_{\delta_0}$, namely that they could be geometrically computed easily.

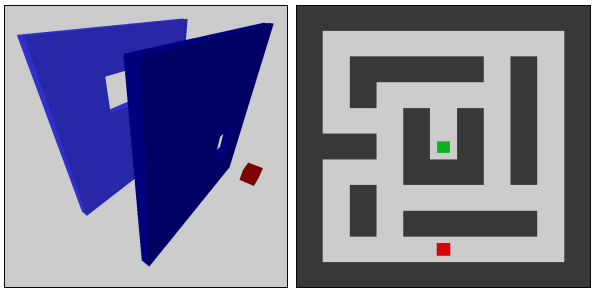

Fig. 6. The environments with obstacles: Barriers (Left), and Maze (Right). Here, the robot translates, but does not rotate.

The parameters of these environments can be found in Figure 7.

Environment	$\lvert\mathcal{C}_{\text{free}}\rvert$	ϵ_{DES}	δ_{n_0}	$L^*_{\delta_0}$
Maze	11150	0.2	5.0	240.0
Barriers	700	0.2	0.5	16.5
Empty ($2D$)	300	0.2	1.0	20.0
Empty (T^3)	248.0520	0.3	0.7	3.4641

Fig. 7. Parameters for the test environments examined.

For the desired path bound, the iteration limit n_0 was computed. Then, out of 500 trials, the probability of successfully generating a short path through the hyperballs over π_0 is computed. The stopping criterion properly selects n_0 so that $\mathbb{P}_{success}$ is greater than the input threshold $\mathbb{P}_{DES}$, which was set to 0.9. The probability of success for the algorithm over time is given in Figure 8 for the chosen environments, where the formula for $\mathbb{P}_{success}$ is derived from Equation 5 to get

$$\mathbb{P}_{success} \geq \mathbb{P}(\Omega^\rho_n)(1-\chi).$$

The bound becomes loose for shorter paths relative to their clearance, as seen for the Torus environment. Otherwise, the bound performs well, and discrepancies are due to simplifying assumptions made during the analysis. The experiments validate that the stopping criterion accurately stops the methods so as to achieve PNO properties.

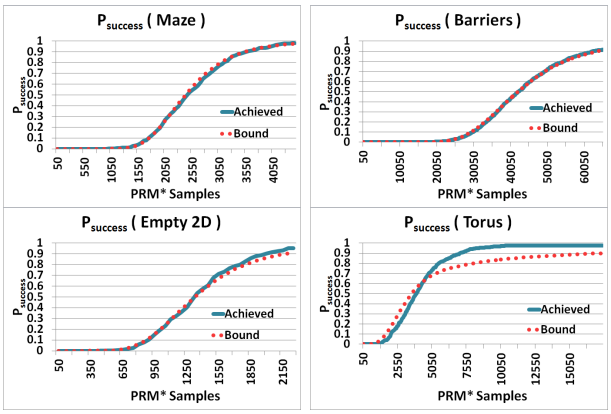

Fig. 8. Probability of returning a path within the bound over time.

VIII. Discussion

This work extends "probabilistic near-optimality" (PNO) properties to PRM*, i.e., an asymptotically optimal but relatively sparse sampling-based roadmap method, and removes dependence on Monte Carlo simulations compared to previous work [20]. The analysis provides tight bounds for path quality, which are validated through simulations.

Due to the difficulty of the motion planning problem, PNO properties often require many samples to provide near-optimal paths with high confidence. But they also provide error and confidence bounds for a given budget of iterations. Difficult problems may require exponentially many samples to provide high-quality solutions, even when considering deterministic and quasi-random sampling. Nevertheless, this bound can inform what level of degradation can still be achieved.

An important future step is the extension of PNO properties to RRT*. Furthermore, extending PNO to sparse methods [25], which add few nodes and reduce computational requirements, is also interesting. Path length bounds are drawn under the assumption of a Euclidean distance metric. For many robotic systems, however, this is not appropriate. Future work should address non-additive cost functions, such as clearance and Hausdorff distance, as well as metrics in spaces such as $SE(3)$.

An effort to consider an analysis for lattice-based sampling has revealed many difficulties. It adds a dependence on the position of the optimal path to the analysis, and involved exponentially many points to cover the space. Furthermore, appropriate lattices must be identified in many dimensions.

This work also only considers only those paths that are within a radius of the optimum one. The provided bound can benefit both from considering other classes of paths as well as using tighter, numerical approximations.

References

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

[2] J. Canny, "The Complexity of Robot Motion Planning," Ph.D. dissertation, MIT, Cambridge, MA, 1988.

[3] H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of Robot Motion: Theory, Algorithms, and Implementations*. Boston, MA: MIT Press, 2005.

[4] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.

[5] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars, "Probabilistic Roadmaps for Path Planning in High-Dimensional Configuration Spaces," *IEEE TRA*, vol. 12, no. 4, 1996.

[6] S. M. LaValle and J. J. Kuffner, "Randomized Kinodynamic Planning," *IJRR*, vol. 20, pp. 378–400, May 2001.

[7] D. Hsu, R. Kindel, J.-C. Latombe, and S. Rock, "Randomized Kinodynamic Motion Planning with Moving Obstacles," *IJRR*, vol. 21, no. 3, pp. 233–255, Mar. 2002.

[8] N. M. Amato, O. B. Bayazit, L. K. Dale, C. Jones, and D. Vallejo, "OBPRM: An Obstacle-based PRM for 3D Workspaces," in *WAFR*, 1998, pp. 155–168.

[9] L. J. Guibas, C. Holleman, and L. E. Kavraki, "A Probabilistic Roadmap Planner for Flexible Objects with a Workspace Medial-Axis-Based Sampling Approach," in *IROS*, October 1999.

[10] B. Raveh, A. Enosh, and D. Halperin, "A Little More, a Lot Better: Improving Path Quality by a Path-Merging Algorithm," *IEEE TRO*, vol. 27, no. 2, pp. 365–370, 2011.

[11] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe, "Analysis of Probabilistic Roadmaps for Path Planning," *IEEE TRA*, vol. 14, no. 1, pp. 166–171, 1998.

[12] A. M. Ladd and L. E. Kavraki, "Measure Theoretic Analysis of Probabilistic Path Planning," *IEEE TRA*, vol. 20, no. 2, pp. 229–242, April 2004.

[13] Z. McCarthy, T. Bretl, and S. Hutchinson, "Proving path non-existence using sampling and alpha shapes," in *ICRA*, May 2012, pp. 2563 – 2569.

[14] S. Karaman and E. Frazzoli, "Sampling-based Algorithms for Optimal Motion Planning," *IJRR*, vol. 30, no. 7, pp. 846–894, June 2011.

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

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

[17] J. Marble and K. E. Bekris, "Asymptotically Near-Optimal Planning with Probabilistic Roadmap Spanners," *IEEE TRO*, vol. 29, pp. 432–444, 2013.

[18] O. Salzman and D. Halperin, "Asymptotically Near-Optimal RRT for Fast, High-quality Motion Planning." Hong Kong, China: ICRA, June 2014.

[19] W. Wang, D. Balkcom, and A. Chakrabarti, "A Fast Online Spanner for Roadmap Construction," *International Journal of Robotics Research*, 2015.

[20] A. Dobson and K. E. Bekris, "A Study on the Finite-Time Near-Optimality Properties of Sampling-Based Motion Planners." Tokyo Big Sight, Tokyo, Japan: IROS, November 2013.

[21] L. Janson and M. Pavone, "Fast Marching Trees: a Fast Marching Sampling-Based Method for Optimal Motion Planning in Many Dimensions." ISRR, December 2013.

[22] K. Hauser, "Randomized Multi-Modal Motion Planning for a Humanoid Robot Manipulation Task," *IJRR*, vol. 30, no. 6, pp. 678–698, 2011.

[23] A. Dobson, G. Moustakides, and K. E. Bekris, "Sampling-Based Roadmap Planners are Probably Near-Optimal after Finite Computation," arXiv, Tech. Rep., April 2014.

[24] A. Kimmel, A. Dobson, Z. Littlefield, A. Krontiris, J. Marble, and K. E. Bekris, "PRACSYS: An Extensible Architecture for Composing Motion Controllers and Planners," in *SIMPAR*, Tsukuba, Japan, 11/2012 2012.

[25] A. Dobson and K. E. Bekris, "Sparse Roadmap Spanners for Asymptotically Near-Optimal Motion Planning," *IJRR*, vol. 33, January 2014.