On the Probabilistic Foundations of Probabilistic Roadmaps (Extended Abstract)

David Hsu[1], Jean-Claude Latombe[2], and Hanna Kurniawati[1]

[1] Department of Computer Science, National University of Singapore
{dyhsu,hannakur}@comp.nus.edu.sg
[2] Department of Computer Science, Stanford University
latombe@cs.stanford.edu.edu

Probabilistic roadmap (PRM) planners [5, 16] solve apparently difficult motion planning problems where the robot's configuration space $\mathcal{C}$ has dimensionality six or more, and the geometry of the robot and the obstacles is described by hundreds of thousands of triangles. While an algebraic planner would be overwhelmed by the high cost of computing an exact representation of the free space $\mathcal{F}$, defined as the collision-free subset of $\mathcal{C}$, a PRM planner builds only an extremely simplified representation of $\mathcal{F}$, called a *probabilistic roadmap*. This roadmap is a graph, whose nodes are configurations sampled from $\mathcal{F}$ with a suitable probability measure and whose edges are simple collision-free paths, *e.g.*, straight-line segments, between the sampled configurations. PRM planners work surprisingly well in practice, but why?

Previous work has partially addressed this question by identifying and formalizing properties of $\mathcal{F}$ that guarantee good performance for a PRM planner using the uniform sampling measure (*e.g.*, [12, 14, 15, 18, 23]). Several systematic experimental studies have also compared various PRM planners, in terms of their sampling and connection strategies (*e.g.*, [7, 8, 21]). However, the underlying question "Why are PRM planners probabilistic?" has received little attention so far, and consequently the importance of probabilistic sampling measures for PRM planning remains poorly understood. Since no inherent randomness or uncertainty exists in the classic formulation of motion planning problems, one may wonder why probabilistic sampling helps to solve them.

Our work attempts to fill this gap by establishing the probabilistic foundations of PRM planning—an effort that that, surprisingly, has not been undertaken before—and re-examining previous work in this context. A full version of this paper will soon appear [11]. The main questions addressed in this work are summarized below:

Why is PRM planning probabilistic? A foundational choice in PRM planning is to avoid the prohibitive cost of computing an exact representation of $\mathcal{F}$. Instead, a PRM planner uses fast probes to test whether sampled configurations and paths are collision-free. So, it never knows the exact shape of $\mathcal{F}$, nor its connectivity. It works very much like a robot exploring an *unknown* environment to build a map. At any moment during planning, many hypotheses on the shape of $\mathcal{F}$ are consistent with the information gathered so far. The probability measure for sampling $\mathcal{F}$ reflects this uncertainty. Hence, PRM planning trades the cost of computing $\mathcal{F}$ exactly against the cost of dealing with uncertainty. This choice is beneficial only if probabilistic sampling is likely to lead to a roadmap that is much smaller in size than that of an exact representation of $\mathcal{F}$ and still represents $\mathcal{F}$ well enough to answer motion planning queries correctly. Note the analogy with PAC learning, where one can expect to learn a concept from examples only if the concept is assumed to have a simple representation.

Why does PRM planning work well? One can think of the nodes of a roadmap as a network of guards watching over $\mathcal{F}$. To guarantee that a PRM planner converges quickly, $\mathcal{F}$ should satisfy favorable "visibility" properties, more specifically a property called *expansiveness* [12]. Perhaps the main contribution of PRM planning has been to reveal, through its emprical success, that many free spaces encountered in practice satisfy this property, despite their high algebraic complexity. This fact was a priori unsuspected, but in retrospect it is not so surprising. Poor visibility is caused by narrow passages, which are unstable geometric features: small random perturbations of the workspace geometry are likely to either eliminate them or make them wider [4]. So, narrow passages rarely occur by accident. Since visibility properties are defined in terms of volume ratios over certain subsets of $\mathcal{F}$, they do not directly depend on $\dim(\mathcal{C})$, the dimensionality of $\mathcal{C}$. This explains why PRM planning scales up reasonably well when $\dim(\mathcal{C})$ increases.

How important is the sampling measure? In every PRM planner, a probability measure prescribes how sampled configurations are distributed over $\mathcal{F}$. Since visibility properties are usually not uniformly favorable across $\mathcal{F}$, non-uniform measures, which strive to identify regions with inferior visibility properties and allocate a higher density of samples to them, have a critical impact on the efficiency of PRM planning. Existing PRM planners use a variety of techniques to localize regions of $\mathcal{F}$ where visibility is expected to be less favorable. Some identify narrow passages in the robot's workspace and map them into configuration space [6, 9, 17, 24, 25], Others, like Gaussian sampling [2] and the bridge test [10], over-sample $\mathcal{C}$, but quickly reject many unpromising samples by detecting local geomeric features suggesting good or poor visibility. Others exploit information gained during roadmap construction to generate and adapt sampling measures [1, 3, 12, 13, 16, 20, 22] Experiments show that the resulting non-uniform sampling measures dramatically improve the performance of PRM planning.

How important is the sampling source? To sample a configuration, a PRM planner needs both a probability measure and a source S of random or deterministic numbers. The "sampling measure", a notion firmly rooted in probability theory, and the "sampling source" are very distinct concepts, but they have often been blurred in the literature. With the use of deterministic sources in PRM planners [19], this distinction becomes important. Typically, a PRM planner uses S to sample a point uniformly from the unit hypercube $[0,1]^{\dim(\mathcal{C})}$ and then maps this point into $\mathcal{C}$ according to the probability measure. The source most commonly used in existing PRM planners is the pseudo-random source that closely approximate the statistical properties of true random numbers. But some deterministic sources can spread samples over $\mathcal{C}$ more evenly by minimizing discrepancy or dispersion [19]. However, experiments show that the choice of the source has limited effect on the efficiency of PRM planning. When $\dim(\mathcal{C})$ is small, low-discrepancy/dispersion deterministic sources achieve some speedup over pseudo-random sources, but this speedup is very modest compared to that achieved by good sampling measures. It also fades away quickly, as $\dim(\mathcal{C})$ increases.

Acknowledgments: D. Hsu's research is supported by NUS grant R252-000-243-112. J.C. Latombe's research is supported by NSF grants ACI-0205671, IIS-0412884, and DMS-0443939, and NIH grant 5R33 LM007295.

References

1. M. Akinc, K.E. Bekris, B. Y. Chen, A.M. Ladd, E. Plaku, and L.E. Kavraki. Probabilistic roadmaps of trees for parallel computation of multiple query roadmaps. In M. Erdmann et al., editors, *Algorithmic Foundations of Robotics VI*, pages 80–89. Springer-Verlag, 2004.
2. V. Boor, M.H. Overmars, and F. van der Stappen. The Gaussian sampling strategy for probabilistic roadmap planners. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 1018–1023, 1999.
3. B. Burns and O. Brock. Toward optimal configuration space sampling. In *Proc. Robotics: Science and Systems*, 2005.
4. S. Chaudhuri and V. Koltun. Smoothed analysis of probabilistic roadmaps. Unpublished manuscript, 2006.
5. 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*, chapter 7. The MIT Press, 2005.
6. M. Foskey, M. Gerber, M.C. Lin, and D. Manocha. A Voronoi-based hybrid motion planner. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 55–60, 2001.
7. R. Geraerts and M.H. Overmars. A comparative study of probabilistic roadmap planners. In J.D. Boissonnat et al., editors, *Algorithmic Foundations of Robotics V*, pages 43–59. Springer, 2002.
8. R. Geraerts and M.H. Overmars. Reachability analysis of sampling based planners. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 406–412, 2005.
9. L. Guibas, C. Holleman, and L.E. Kavraki. A probabilistic roadmap planner for flexible objects with a workspace medial-axis based sampling approach. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 254–260, 1999.
10. D. Hsu, T. Jiang, J. Reif, and Z. Sun. The bridge test for sampling narrow passages with probabilistic roadmap planners. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 4420–4426, 2003.
11. D. Hsu, J.C. Latombe, H. Kurniawati. On the Probabilistic Foundations of Probabilistic Roadmaps. To appear in *Int. J. of Robotics Research*, 2006.

12. D. Hsu, J.C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 2719–2726, 1997.
13. D. Hsu, G. Sanchez-Ante, and Z. Sun. Hybrid PRM sampling with a cost-sensitive adaptive strategy. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 3885–3891, 2005.
14. L.E. Kavraki, M.N. Kolountzakis, and J.C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Trans. on Robotics & Automation*, 14(1):166–171, 1998.
15. L.E. Kavraki, J.C. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot path planning. In *Proc. ACM Symp. on Theory of Computing*, pages 353–362, 1995.
16. L.E. Kavraki, P. Švestka, J.C. Latombe, and M.H. Overmars. Probabilistic roadmaps for path planning in highdimensional configuration space. *IEEE Trans. on Robotics & Automation*, 12(4):566–580, 1996.
17. H. Kurniawati and D. Hsu. Workspace importance sampling for probabilistic roadmap planning. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 1618–1623, 2004.
18. A.M. Ladd and L.E. Kavraki. Theoretic analysis of probabilistic path planning. *IEEE Trans. on Robotics & Automation*, 20(2):229–242, 2004.
19. S.M. LaValle, M.S. Branicky, and S.R. Lindemann. On the relationship between classical grid search and probabilistic roadmaps. *Int. J. Robotics Research*, 23(7/8):673–692, 2004.
20. S.M. LaValle and J.J. Kuffner. Randomized kinodynamic planning. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 473–479, 1999.
21. M.A. Morales A., R. Pearce, and N.M. Amato. Metrics for analyzing the evolution of c-space models. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2006.
22. M.A. Morales A., L. Tapia, R. Pearce, S. Rodriguez, and N.M. Amato. A machine learning approach for featuresensitive motion planning. In M. Erdmann et al., editors, *Algorithmic Foundations of Robotics VI*, pages 361–276. Springer-Verlag, 2004.
23. P. Švestka. *On probabilistic completeness and expected complexity for probabilistic path planning.* Tech. Rep. UU-CS-1996-08, Utrecht University, Dept. of Information & Computing Sciences, Utrecht, The Netherlands, 1996.
24. J.P. van den Berg and M.H. Overmars. Using workspace information as a guide to non-uniform sampling in probabilistic roadmap planners. *Int. J. Robotics Research*, 24(12):1055–1071, 2005.
25. Y. Yang and O. Brock. Adapting the sampling distribution in PRM planners based on an approximated medial axis. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2004.