On the Probabilistic Foundations of Probabilistic Roadmap Planning

David Hsu* **Jean-Claude Latombe+** **Hanna Kurniawati***

*Department of Computer Science, National University of Singapore, Singapore

+Department of Computer Science, Stanford University, Stanford, CA, USA

Abstract. Why are probabilistic roadmap (PRM) planners "probabilistic"? This paper establishes the probabilistic foundations of PRM planning – something that, surprisingly, has not been done before – and re-examines previous work in this context. It shows that the success of PRM planning depends mainly and critically on the assumption that the configuration space *C* of a robot often verifies favorable "visibility" properties not directly dependent on the dimensionality of *C*. A promising way of speeding up PRM planners is to extract partial knowledge on such properties during roadmap construction and use it to adjust the sampling measure continuously.

1. Introduction

Probabilistic roadmap (PRM) planners [CLH+05] solve seemingly difficult motion planning problems such as the one in Figure 1, where the robot's configuration space *C* is 6-D and the environment consists of tens of thousands of triangles. While an algebraic planner would be overwhelmed by the high cost of computing an exact representation of the free space *F*, defined as the collision-free subset of *C*, a PRM planner builds only an extremely simplified representation of *F*, called a probabilistic roadmap. The nodes of the roadmap *R* are configurations sampled from *F* with a suitable probability measure. The edges of *R* are simple collision-free paths, *e.g.*, straight-line segments, between the sampled configurations. PRM planners work surprisingly well in practice. Why?

This question leads us to establish the probabilistic foundations of PRM planning – something that, surprisingly, has not been done before – and re-examine previous work in this context. We emphasize the distinction between a sampling *measure*, a notion firmly rooted in the probability theory, and a sampling *source*, with the intent to identify the key components determining the efficiency of PRM planning. The main issues addressed in this paper are summarized below:

Figure 1: A practical motion planning problem.

- **Why is PRM planning "probabilistic"?** A foundational choice in PRM planning is to avoid computing an exact representation of *F*. So the planner never knows the exact shape of *F*. At any moment during planning, many hypotheses on *F* are consistent with the configurations sampled so far. The probability measure for sampling *F* reflects this uncertainty. Hence, PRM planning trades the cost of computing *F* exactly against the cost of dealing with uncertainty. This choice is beneficial only if a small roadmap can represent the shape of *F* well enough to answer motion-planning queries correctly.

- **Why does PRM planning work well?** One can think of the nodes of a roadmap as a network of guards *watching* over *F*. To guarantee that a PRM planner converges quickly, *F* should satisfy favorable "visibility" properties. Perhaps the main contribution of PRM planning is to reveal that many free spaces encountered in practice satisfy such properties, despite their high algebraic complexity. Since visibility properties can be defined in terms of volume ratios over certain subsets of *F*, they do not directly depend on dim(*C*), the dimensionality of *C*. This explains why PRM planning scales up reasonably well when dim(*C*) increases.
- **How important is the sampling measure?** In every PRM planner, a probability measure prescribes how sampled configurations are distributed over *F*. Since visibility properties are usually not uniformly favorable over *F*, this measure plays a critical role in the efficiency of PRM planning by allocating a higher density of samples to regions with poor visibility properties. Existing PRM planners use mostly simple, heuristic estimates of visibility properties, but experiments show that they dramatically improve the performance of PRM planning.
- **How important is the sampling source?** A PRM planner also needs a source *S* of *uniformly* distributed pseudo-random or deterministic numbers for sampling *C*. For example, a planner may use *S* to pick a point uniformly from $[0,1]^{\dim(C)}$ and then maps the point into *C* according to a given probability measure. The source *S* has only a limited effect on the efficiency of PRM planning. When dim(*C*) is small, low-discrepancy or low-dispersion deterministic sources achieve some speedup over pseudo-random sources [LBL04]. The speedup is, however, very modest compared to that achieved by good sampling measures. It also fades away quickly, as dim(*C*) increases.

This paper does not introduce any new PRM planner or sampling strategy. Instead, its contribution is to articulate a coherent framework centered on the probabilistic foundations of PRM planning and evaluate several ideas, considered separately before, in this framework. It brings new understanding of what makes PRM planning effective, which in turn may help us to design better planners in the future.

2. Why is PRM planning "probabilistic"?

For many robots, computing an exact representation of the free space *F* takes prohibitive time, but fast, exact algorithms exist to test whether a given configuration or path is collision-free. PRM planners use two *probes* based on such algorithms to access geometric information from the configuration space *C*:

- For any $\boldsymbol{q} \in C$, `FreeConf`($\boldsymbol{q}$) is true if and only if $\boldsymbol{q} \in F$.
- For any pair $\boldsymbol{q}, \boldsymbol{q}' \in C$, `FreePath`($\boldsymbol{q},\boldsymbol{q}'$) is true if and only if $\boldsymbol{q}$ and $\boldsymbol{q}'$ can be connected with a straight-line path lying entirely in *F*.

The choice of using only these two probes is foundational for PRM planning. Since a PRM planner does not compute the exact shape of F, it never gains this information. So, it acts just like a robot building a map of an *unknown* environment. At any moment, many hypotheses on F are consistent with the information gathered so far by the probes, and each hypothesis has some probability of being correct. The probabilistic nature of PRM planners comes from the fact that this uncertainty is modeled implicitly by a probability measure over the hypotheses.

In this paper, we use the following scheme, which we call `Basic-PRM`, as a reference planner. Like the original PRM planner [KSLO96], it operates in two stages, roadmap construction and query.

- **Roadmap construction.** The procedure below takes a single input argument N, the number of nodes of the roadmap R to be constructed. The nodes of R are collision-free configurations sampled from F. The edges represent collision-free straight-line paths between the nodes.

Procedure `Roadmap-Construction`(N)

1. **repeat** until N nodes have been generated
2.   Sample a configuration $\boldsymbol{q}$ from C uniformly at random.
3.   **if** `FreeConf`($\boldsymbol{q}$) is true **then** add $\boldsymbol{q}$ as a new node of R.
4.   **for** every node $\boldsymbol{q}'$ of R such that $\boldsymbol{q}' \neq \boldsymbol{q}$ **do**
5.   **if** `FreePath`($\boldsymbol{q}$, $\boldsymbol{q}'$) is true **then** add ($\boldsymbol{q}$, $\boldsymbol{q}'$) as a new edge of R.
6. **return** R.

Most PRM planners use better sampling strategies than the uniform random one in Line 2, as well as better connection strategies in Lines 4–5.

A sampling strategy (π, S) is characterized by a probability *measure* π that prescribes how sampled configurations are distributed over C and a *source* S of uniformly distributed pseudo-random or deterministic numbers. We will show in Sections 4–5 that designing good sampling measures is one of the most promising ways to speed up PRM planning.

- **Roadmap query.** A query is defined by two configurations $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$ in F. Given a roadmap R, the procedure `Roadmap-Query` tries to connect each $\boldsymbol{q}_i$, i=1,2, to a corresponding node $\boldsymbol{v}_i$ of R. For each $\boldsymbol{q}_i$, it samples uniformly at random K configurations so that for each such configuration $\boldsymbol{q}$, `FreePath`($\boldsymbol{q}_i$,$\boldsymbol{q}$) is true. It then checks whether there is a node $\boldsymbol{v}_i$ of R such that `FreePath`($\boldsymbol{q}$,$\boldsymbol{v}_i$) is true. If either $\boldsymbol{q}_1$ or $\boldsymbol{q}_2$ cannot be connected to a node of R, `Roadmap-Query` returns FAILURE. Otherwise, it searches for a path in R between $\boldsymbol{v}_1$ and $\boldsymbol{v}_2$. If one is found, it returns a path between $\boldsymbol{q}_1$ and $\boldsymbol{q}_2$. Otherwise, it returns NO PATH.

If `Roadmap-Query` returns a path, the answer is always correct, but the NO PATH answer may not be correct, as disconnected components of R may lie in the same connected component of F. The answer FAILURE means that R is insufficient to answer the query.

3. Why does PRM planning work well?

In general, `Basic-PRM` may return incorrect NO PATH or FAILURE answers, but the efficiency of PRM planners in practice indicates that these outcomes have low probability γ. Experiments show that even in complex geometric environments, γ often converges to 0 quickly, as N, the number of roadmap nodes, increases (see, *e.g.*, Figure 2). Yet one can also easily construct apparently simple environments where PRM planners perform terribly (Figure 3). Together these two examples suggest that many environments encountered in practice satisfy favorable properties that PRM planners exploit well. What are these properties?

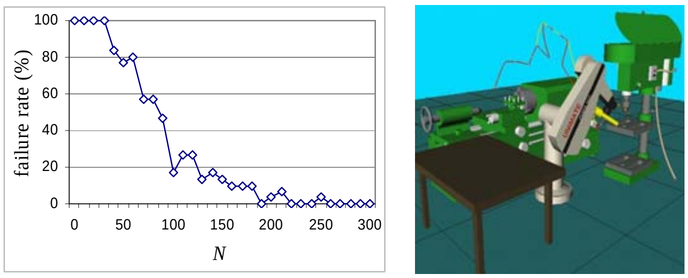

Figure 2: The experimental convergence rate of `Basic-PRM`. The plot shows the percentage of unsuccessful outcomes out of 100 independent runs for the same query, as the number of roadmap nodes increases.

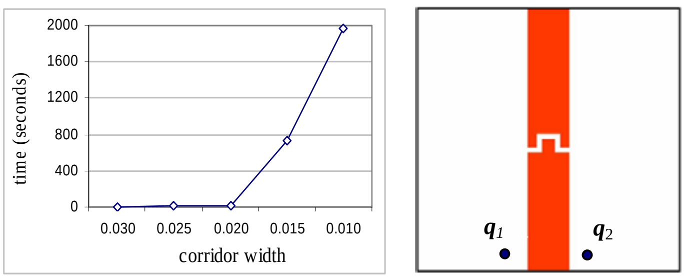

Figure 3: A difficult example for PRM planning. F consists of two rectangular chambers connected by a narrow corridor. The plot shows the average running time for `Basic-PRM` to connect the two query configurations, as the corridor width decreases.

We now review results from [KLMR95, HLM97], showing that if F satisfies a rather general visibility property, called *expansiveness*, then `Basic-PRM` answers planning queries correctly with high probability. In the following, the phrase "with high (low) probability in n" means that the probability converges to 1 (0) at an exponential rate, as n increases.

3.1. Visibility in the free space

We say that two points $\boldsymbol{q}$ and $\boldsymbol{q}'$ in F *see* each other if `FreePath`($\boldsymbol{q}$, $\boldsymbol{q}'$) is true. The *visibility set* of $\boldsymbol{q} \in F$ is the set $V(\boldsymbol{q}) = \{ \boldsymbol{q}' \in F \mid$ `FreePath`$(\boldsymbol{q}, \boldsymbol{q}')$ is true$\}$. The definition of visibility set is extended to any set M of points in F by setting $V(M) = \cup_{q \in M} V(\boldsymbol{q})$.

Definition 1. Given a constant $\varepsilon \in (0,1]$, a point $\boldsymbol{q} \in F$ is ε-*good* if it sees at least an ε-fraction of F, *i.e.*, if $\mu(V(\boldsymbol{q})) \geq \varepsilon \times \mu(F)$, where $\mu(S)$ denotes the volume of S for any $S \subseteq C$. F is ε-*good* if every point $\boldsymbol{q} \in F$ is ε-good.

Definition 2. A roadmap R provides *adequate coverage* of an ε-good free space F if the subset of F not seen by any node of R has volume at most $(\varepsilon/2)\mu(F)$.

Theorem 1 [KLMR95]**.** If F is ε-good, then with high probability in N, `Roadmap-Construction` generates a roadmap that provides adequate coverage of F.

Theorem 2 [KLMR95]**.** If a roadmap provides adequate coverage of F, then `Roadmap-Query` returns FAILURE with low probability in K.

Adequate coverage only protects us from FAILURE, but does not prevent an incorrect NO PATH answer, because ε-goodness is too weak to imply anything on roadmap connectivity. A stronger property is needed to "link" a visibility set to its complement in F.

Definition 3. Let F' be a connected component of F, G be any subset of F', and β be a number in (0,1]. The β-LOOKOUT of G is the set of all points in G that see at least a β-fraction of the complement of G in F':

$$\beta\text{-LOOKOUT}(G) = \{\boldsymbol{q} \in G \mid \mu(V(\boldsymbol{q})\backslash G) \geq \beta\times\mu(F' \backslash G)\}.$$

Suppose that the volume of β-LOOKOUT(G) is at least $\alpha\times\mu(G)$. If either α or β is small, then it would be difficult to sample a point in G *and* another in $F'\backslash G$ so that the two points see each other, hence to build a roadmap that represents the connectivity of F' well. This happens in the free space of Figure 3 when the corridor is very narrow. These considerations lead to the concept of expansiveness.

Definition 4: Let ε, α, and β be constants in (0,1]. A connected component F' of F is $(\varepsilon,\alpha,\beta)$*-expansive* if (i) every point $\boldsymbol{q}\in F'$ is ε*-good* and (ii) for any set M of points in F', $\mu(\beta\text{-LOOKOUT}(\boldsymbol{V}(M))) \geq \alpha\times\mu(\boldsymbol{V}(M))$. F is $(\varepsilon,\alpha,\beta)$*-expansive*, if its connected components are all $(\varepsilon,\alpha,\beta)$-expansive.

Theorem 3 [HLM97]: If F is $(\varepsilon,\alpha,\beta)$-expansive, then with high probability in N, `Roadmap-Construction` generates a roadmap whose connected components have one-to-one correspondence with those of F.

Expansiveness guarantees that the visibility set $\boldsymbol{V}(M)$ of any set M of points in a connected component F' of F has a large lookout. So it is easy to sample at random a set of configurations and construct a roadmap that provides good coverage of F and represent the connectivity of F well. The values of ε, α, and β measure the extent to which F is expansive. For example, if F is convex, then $\varepsilon=\alpha=\beta=1$. The larger these values are, the smaller N needs to be for `Basic-PRM` to answer queries correctly. Although for a given motion planning problem, we often cannot compute these values in advance, they characterize the nature of free spaces in which PRM planning works well.

3.2. What does the empirical success of PRM planners imply?

In practice, a small number of roadmap nodes are often sufficient to answer queries correctly. This frequent success suggests that the main reason for the empirical success of PRM planners is that free spaces encountered in practice often satisfy favorable visibility properties, such as expansiveness. PRM planners scale up well when dim(C) increases, because visibilities properties can be defined in terms of volume ratios over subsets of F and do not directly depend on dim(C). So, the empirical success of PRM planning says as much about the nature of motion-planning problems encountered in practice as about the algorithmic efficiency of PRM planning. The fact that many free spaces, despite their high algebraic complexity, verify favorable visibility properties is not obvious a priori. An important contribution of PRM planning is to reveal this fact.

We have no proof that expansiveness is the *minimal* property that F must satisfy for PRM planners to work well, but few alternatives exist (*e.g.*, path clearance and Lipschitz ε-RB) and they are more specific. However, since the values of ε, α, and β are determined by the worst configurations and lookouts in F, they do not reflect the variation of visibility properties over F. This is precisely what non-uniform sampling measures described below try to exploit.

4. How important is the sampling measure?

Most PRM planners employ non-uniform sampling measures that dramatically improve performance. To illustrate, Figure 4 compares the average running times of three versions of `Basic-PRM` using sampling strategies with different measures**:** the uniform strategy, the two-phase connectivity expansion strategy [KSLO96], and the Gaussian strategy [BOvdS99]. The last two strategies perform much better than the uniform one. How can such improvement be explained? What information can a PRM planner use to bias the sampling measure to its advantage?

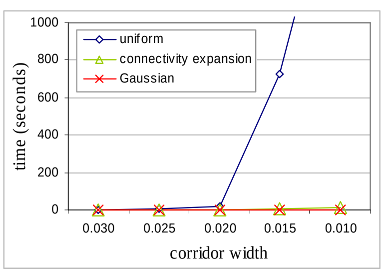

Figure 4: Comparison of three strategies with different sampling measures. The plot shows the average running times over 30 runs on the problem in Figure 3, as the corridor width decreases.

If nothing is assumed on F, all hypotheses on the shape of F are equally likely. There is no reason to sample one region of C more densely than another, and the uniform sampling measure is the best that a PRM planner can use. More generally, with no prior assumptions, there is little that we can say about the expected performance of PRM planners. If we persist in using PRM planners, it must be that F is expected to satisfy certain favorable properties. Note here the analogy with the theory of PAC learning, where one can expect to learn a concept from examples only if the concept is assumed to have a simple representation. Similarly, *we can expect a PRM planner to work well – i.e., to "learn" the shape of F from sampled configurations – only if we assume that F satisfies favorable visibility properties, which allow it to be adequately represented by a small roadmap.*

Now, if F is expansive, can non-uniform sampling measures work better than the uniform one? Since visibility properties are not uniformly favorable over F, a PRM planner should exploit the partial knowledge acquired during roadmap construction to identify regions with poor visibility properties and adjust the probability measure to sample these regions more densely. Now not only is the sampling measure non-uniform over F, but also it changes over time. In each sampling operation, the optimal measure is the one that minimizes the expected number of remaining sampling operations needed to reach a good roadmap.

The problem of constructing good sampling measures is still poorly understood. Existing strategies mostly rely on simple, heuristic estimates of visibility properties, for instance:

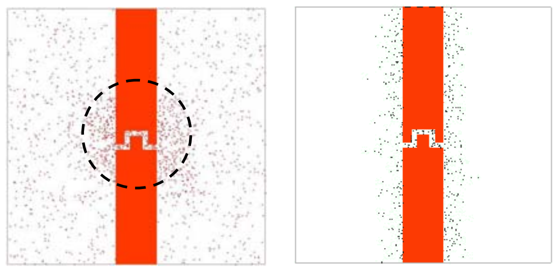

Figure 5: Sampled configurations generated by (*a*) the two-phase connectivity expansion strategy and (*b*) the Gaussian strategy.

- The *two-phase connectivity expansion strategy* [KSLO96] builds an initial roadmap by sampling C uniformly at random. While doing so, it identifies the nodes that frequently fail to connect to other nodes nearby. Then the strategy samples more configurations around these identified nodes. The final distribution of sampled configurations is denser in regions having poor visibility. See the circled region in Figure 5*a*.
- In each sampling operation, the *Gaussian strategy* [BOvdS99] samples a pair of configurations, whose distance between them is chosen according to the Gaussian measure. If exactly one configuration lies in F, this configuration is retained as a roadmap node. Otherwise, both configurations are discarded. This strategy yields a distribution of sampled configurations that is denser near the boundary of F (Figure 5*b*). The intuition is that regions with poor visibility often lie near the boundary.

Figure 4 shows that these two strategies are effective in exploiting the non-uniformity of visibility properties in F. When the corridor width is small, regions near the corridor have poor visibility, and the non-uniform strategies achieve huge speedup over the uniform one. As the corridor width increases, visibility properties become more uniformly favorable. The benefit of non-uniform sampling then decreases.

These two non-uniform strategies are chosen here for illustration purposes only. Others have been proposed (see [CLH+05] for a survey), including one that exploits visibility properties explicitly [SLN00]. Experiments show that they usually achieve *major* speedup over the uniform strategy. Connection strategies (not discussed here for lack of space) also exploit visibility properties in F to improve efficiency.

5. How important is the sampling source?

We have mentioned in Section 2 that a sampling strategy (π, S) is characterized by a probability measure π and a source S. The most commonly used source in PRM planning is the pseudo-random source S_{ran}. Given a fixed seed, S_{ran} generates a deterministic sequence of numbers that approximate closely the statistical properties of true random numbers. In particular, a pseudo-random sequence is slightly irregular to simulate the effect that each number is chosen independently. In the proofs of Theorems 1–3, this independence guarantees that samples spread evenly over F according to the uniform measure. However, deterministic sources can achieve the same goal, sometimes even better [LBL04]. A familiar deterministic source is a *grid*. In this section, we compare pseudo-random and deterministic sources. We also compare the impact of sampling sources and sampling measures on the overall efficiency of PRM planning.

In our experiments, we use a pseudo-random source S_{ran} as well as two deterministic sources, the Halton sequence S_{hal} and the incremental discrepancy-optimal sequence S_{opt}, both of which have been reported to often outperform S_{ran} [GO02, LBL04]. We then pair each source with two probability measures, the uniform measure π_U and the measure π_G used in the Gaussian strategy. This leads to six sampling strategies $\{\pi_U,\pi_G\}\times\{S_{ran},S_{hal},S_{opt}\}$, each embedded in a distinct version of `Basic-PRM` to be tested experimentally.

	π_U	π_G
S_{ran}	1.0	92
S_{hal}	1.4	23
S_{opt}	4.0	92

(*a*)

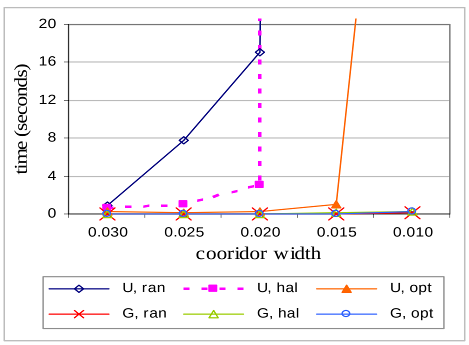

(*b*)

Figure 6: Comparison of six sampling strategies on the problem of Figure 3 when (*a*) the corridor width is set to 0.03 and (*b*) the width decreases.

- **The sampling measure versus the sampling source.** Figure 6*a* compares the six strategies on the example in Figure 3, when the corridor width is set to 0.03. Each table entry gives the ratio of the running time of the uniform random strategy (π_U,S_{ran}) versus that of the strategy of the entry. So, the table reports the *speedup* over (π_U,S_{ran}). The running times for (π_U,S_{ran}) and (π_G,S_{ran}) are averaged over 30 independent runs. The second column (π_U) shows that S_{hal} and S_{opt} indeed achieve some speedup over S_{ran}, but far greater speedup is achieved by switching to π_G. Furthermore, the advantage of S_{hal} and S_{opt} over S_{ran} observed with π_U vanishes when we switch to π_G. These results are reinforced in Figure 6*b*, which plots the running times of the six strategies, as the corridor width decreases. The three curves bundled together at the bottom of the plot all correspond to strategies using π_G, demonstrating the importance of the sampling measure on the overall efficiency of the planner. Similar results have been obtained on more realistic problems, *e.g.*, the one in Figure 8, in which a 6-degrees-of-freedom robot manipulator needs to access the bottom of a car through the narrow space between the lift supports.

	π_U	π_G
S_{ran}	1.0	40.3
S_{hal}	3.9	33.2
S_{opt}	0.9	42.2

Figure 7: Comparison of six sampling strategies on a more realistic problem.

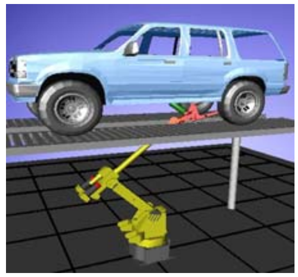

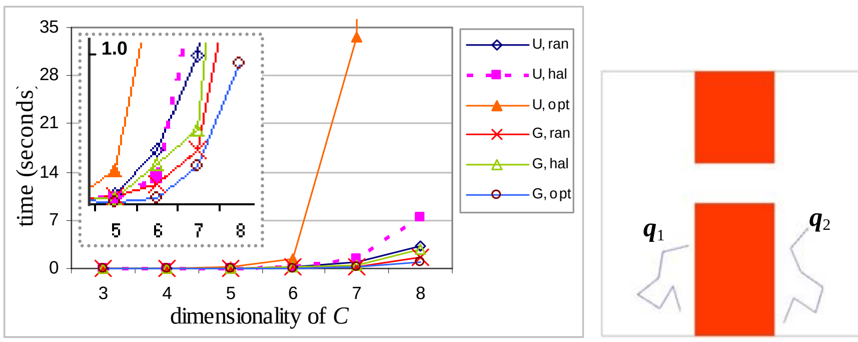

Figure 8: The dependence of six sampling strategies on dimensionality of C. The inset in the right plot is a zoom of the lower portions of the curves.

- **Dependence on dimensionality.** The main theoretical basis for deterministic sources is that that they minimize criteria such as discrepancy or dispersion. Yet the computational cost of maintaining a fixed discrepancy or dispersion increases *exponentially* with dim(C) [Mat99]. When dim(C) $\geq$ 6, roadmaps are necessarily sparse, just like very low resolution grids. Hence, the advantage that deterministic sources can possibly achieve over pseudo-random sources fades away as dim(C) increases. Figure 8 gives an example, showing the running times of the six strategies as dim(C) increase from 3 to 8. The robot is a planar linkage with a mobile base. We increase dim(C) by adding more links. Figure 8 shows that the running time of (π_U,S_{opt}) increases quickly with dim(C). The increase is slower with (π_U,S_{hal}) and even slower with (π_U,S_{ran}). It is interesting to observe that (π_U,S_{hal}) performs slightly better than (π_U,S_{ran}) when dim(C) $\leq$ 6, but worsens afterwards (see the inset in the plot). The three strategies using π_G all have only moderate increases in running times. As dim(C) increases, visibility properties become less uniformly favorable over F, and the advantage of π_G over π_U grows.

For lack of space, we cannot present all our experimental results or the experimental setup, but all the results show that the sampling measure is far more important than the sampling source in determining the efficiency of PRM planners.

6. Conclusion

The success of PRM planning depends mainly and critically on the assumption that, in practice, free spaces often verify favorable visibility properties. Non-uniform sampling measures dramatically improve the efficiency of PRM planning by exploiting these properties. In contrast, the choice of sampling sources has only small impact.

This paper suggests that a promising way to speed up PRM planning is to design better sampling strategies (and perhaps connection strategies as well) by exploiting the partial knowledge acquired during roadmap construction to adjust the sampling measure continuously. Initial work in this direction has appeared recently [BB05, HSS05]. Another possibility is to design probes `FreeConf` and `FreePath` that return a local description of *F*, rather than only a binary outcome, allowing the planner to better identify regions with poor visibility.

Acknowledgements. This work was done in part while J.C. Latombe was at NUS, supported by the Kwan Im Thong Hood Cho Temple Professorship. J.C. Latombe's research is funded by NSF grants ACI-0205671, IIS-0412884, and DMS-0443939, and NIH grant 5R33 LM007295.

References

[BOvdS99] V. Boor, M.H. Overmars, and A.F. van der Stappen. The Gaussian sampling strategy for probbailistic roadmap planners. *Proc. IEEE Int. Conf. on Robotics & Automation*, pp. 1018-1023, 1999.

[BB05] B. Burns and O. Brock. Sampling-based motion planning using predictive models. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2005.

[CLH+05] 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, MIT Press, 2005.

[GO02] R. Geraerts and M.H. Overmars. A comparative study of probabilistic roadmap planners. In J.D. Boissonnat et al. (eds.), *Algorithmic Foundations of Robotics V*, pp. 43-57, Springer, 2004.

[HLM96] D. Hsu, J.C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pp. 2219-2226, 1997.

[HSS05] D. Hsu, G. Sánchez-Ante, and, Z. Sun. Hybrid PRM sampling with a cost-sensitive adaptive strategy. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2005.

[KLMP95] L.E. Kavraki, J.C. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot motion planning. In *Proc. ACM Symp. on Theory of Computing (STOC)*, pp. 353-362, 1995.

[KSLO96] L.E. Kavraki, P. Svestka, J.C. Latombe, M.H. Overmars, Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. on Robotics & Automation*, 12:566–580, 1996.

[LBL04] S.M. LaValle, M.S. Branicky, and S.R. Lindermann. On the relationship between classical grid search and probabilistic roadmaps. *Int. J. of Robotics Research*, 23(7-8):673-692, 2004.

[Mat99] J. Matousek. *Geometric Discrepancy*. Springer-Verlag, Berlin, 1999.

[SLN00] T. Siméon, J.P. Laumond, and C. Nissoux. Visibility-based probabilistic roadmaps for motion planning. *J. of Advanced Robotics*, 14(6):477-494, 2000.