# On the Probabilistic Foundations of Probabilistic Roadmap Planning

David Hsu[1, 2], Jean-Claude Latombe[3, *], and Hanna Kurniawati[1]

[1]Department of Computer Science, National University of Singapore, Singapore
[2]Singapore MIT Alliance, Singapore
[3]Department of Computer Science, Stanford University, Stanford, CA, USA

**Summary.** Why are probabilistic roadmap (PRM) planners "probabilistic"? This paper tries to establish the probabilistic foundations of PRM planning 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 that are 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 this knowledge to adjust the sampling measure continuously. This paper also shows that the choice of the sampling source—pseudo-random or deterministic—has small impact on a PRM planner's performance, compared to that of the sampling measure. These conclusions are supported by both theoretical arguments and empirical results.

## 1 Introduction

Probabilistic roadmap (PRM) planners [3, Chapter 7] solve seemingly difficult motion planning problems such as the one in Fig. 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 a 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?

Previous work has partially addressed this question by identifying and formalizing free space properties that provide sufficient conditions to guarantee that a PRM planner using a uniform sampling measure works well. However, the underlying question "Why are PRM planners probabilistic?" has received little attention so far, and consequently the role of non-uniform sampling measures in

[*] Part of this work was completed while the author was at the National University of Singapore, supported by the Kwan Im Thong Hood Cho Temple Professorship.

PRM planning remains poorly understood. Since no inherent randomness or uncertainty exists in the classic formulation of motion planning problems like the one depicted in Fig. 1, one may wonder why probabilistic sampling helps to solve them.

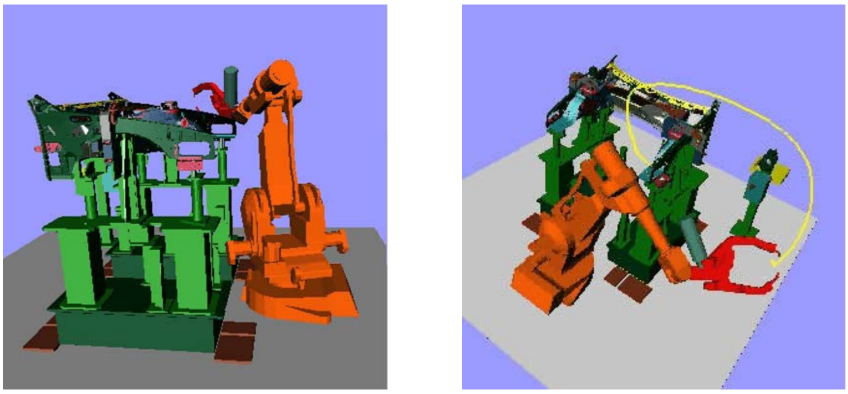

**Fig. 1.** A practical motion planning problem.

In this paper, we attempt to fill this gap, with the intent of identifying promising directions to improve future PRM planners. We introduce the probabilistic foundations of PRM planning (Section 2). We then examine previous work in this context and argue that the empirical success of PRM planning tells us as much about the nature of motion planning problems encountered in practice as about PRM planning itself (Section 3). We emphasize the important distinction between the sampling *measure*, a notion firmly rooted in probability theory, and the sampling *source*, and show that the source has small impact on a planner's performance compared to the measure (Sections 4 and 5).

The main questions addressed in this paper are summarized below:

- **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*, in particular, 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 *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 finds a solution quickly whenever one exists, *F* should satisfy favorable "visibility" properties. A key contribution of PRM planning is to reveal that in practice, many free spaces 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 in general 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 needs a source *S* of *uniformly* distributed pseudo-random or deterministic numbers for sampling *C*. Usually, it calls *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 [13]; however, the speedup is very modest compared to that achieved by good sampling measures and 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 [16]. 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. 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 set of hypotheses.

In this paper, we use the following scheme, which we call `Basic-PRM`, as a reference planner. Like the original PRM planner [12], it operates in two stages, roadmap construction and roadmap query.

- **Roadmap construction.** The procedure below takes a single input argument $N$, the number of nodes for 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 q from C uniformly at random.
3.    if FreeConf(q) is true then add q as a new node of R.
4.    for every node q' of R such that q' ≠ q do
5.       if FreePath(q, q') is true then add (q, 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 ($\pi$, $S$) is characterized by a probability *measure* $\pi$ 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 node 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 so, $\boldsymbol{q}_i$ and $\boldsymbol{v}_i$ can be connected via $\boldsymbol{q}$. 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.

Let us now return to the question "Why is PRM planning probabilistic?". Suppose that while constructing a roadmap, `Roadmap-Construction` could maintain a representation ($H$,$\eta$), where $H$ is the set of all hypotheses over the shape of $F$ and $\eta$ is a probability measure that assigns to each hypothesis in $H$ the probability of it being correct. Suppose further that we can define what a good roadmap is (see Section 3). Then, in each iteration of `Roadmap-`

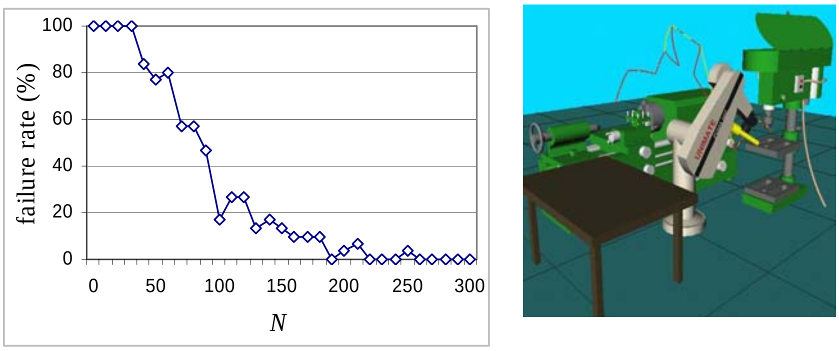


**Fig. 2.** The experimental convergence rate of `Basic-PRM`. The plot shows the percentage of unsuccessful outcomes out of 100 independent runs for the same query in the environment shown on the right, as the number of roadmap nodes increases.

`Construction`, the optimal sampling measure $\pi^*$ is the one that minimizes the expected number of remaining iterations until a good roadmap is reached, and $\pi^*$ can be inferred from $(H,\eta)$. However, maintaining $(H,\eta)$ explicitly is expensive. So existing PRM planners use heuristics to select the sampling measure $\pi$ (see Section 4).

# 3 Why does PRM planning work well?

In general, `Basic-PRM` may return an incorrect NO PATH or FAILURE answer with some probability $\gamma$, but the efficiency of PRM planners in practice indicates that $\gamma$ is usually small. Experiments show that even in complex geometric environments, $\gamma$ often converges to 0 quickly, as $N$, the number of roadmap nodes, increases (Fig. 2). Yet one can also easily construct apparently simple environments where PRM planners perform terribly (Fig. 3). Together, these two examples suggest that many environments encountered in practice satisfy favorable properties that PRM planners exploit well. What are these properties?

We now review results from [9, 11], 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

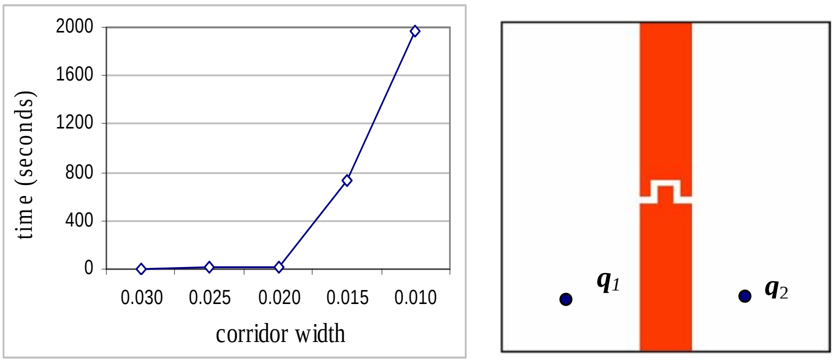


**Fig. 3.** A difficult example for PRM planning. *F* consists of two rectangular chambers connected by a narrow corridor. The plot shows the average running time for `Basic-PRM` to connect the two query configurations, as the corridor width decreases.

definition of a visibility set is extended to any subset *M* of points in *F* by defining $V(M) = \cup_{q \in M} V(\boldsymbol{q})$.

The first two theorems below say that if *F* satisfies a property called $\varepsilon$-goodness, then `Basic-PRM` generates a roadmap that provides good *coverage* of *F* so that `FAILURE` rarely occurs.

**Definition 1.** Given a constant $\varepsilon \in (0,1]$, a point $\boldsymbol{q} \in F$ is *$\varepsilon$-good* if it sees at least an $\varepsilon$-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 *$\varepsilon$-good* if every point $\boldsymbol{q} \in F$ is $\varepsilon$-good.

**Definition 2.** A roadmap *R* provides *adequate coverage* of an $\varepsilon$-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.** If *F* is $\varepsilon$-good, then with high probability in *N*, `Roadmap-Construction`(*N*) generates a roadmap that provides adequate coverage of *F* [11].

**Theorem 2.** If a roadmap provides adequate coverage of *F*, then `Roadmap-Query` returns FAILURE with low probability in *K* [11].

(Recall that *K* is the number of configurations sampled randomly in the neighborhood of each of the query configurations. See Section 2.)

Adequate coverage only protects us from FAILURE, but does not prevent an incorrect NO PATH answer, because $\varepsilon$-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 $\beta$ be a number in (0,1]. The $\beta$-LOOKOUT of *G* is the set of all points in *G* that see at least a $\beta$-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 $\beta$-LOOKOUT($G$) is at least $\alpha\times\mu(G)$. If either $\alpha$ or $\beta$ 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 Fig. 3 when the corridor is very narrow. These considerations lead to the concept of expansiveness.

**Definition 4.** Let $\varepsilon$, $\alpha$, and $\beta$ 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 *$\varepsilon$-good* and (ii) for any set $M$ of points in $F'$, $\mu(\beta\text{-LOOKOUT}(V(M))) \geq \alpha\times\mu(V(M))$. $F$ is *($\varepsilon,\alpha,\beta$)-expansive*, if its connected components are all ($\varepsilon,\alpha,\beta$)-expansive.

**Theorem 3.** 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 $F$ [9].

Expansiveness guarantees that the visibility set $V(M)$ 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 both provides good coverage of $F$ and represents the connectivity of $F$ well. The values of $\varepsilon$, $\alpha$, and $\beta$ 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. If a connected component $F'$ of $F$ had very small values of $\varepsilon$, $\alpha$, and $\beta$, then a planner would likely encounter a set $M$ of sampled nodes such that $V(M)$ has a small lookout. It would then be difficult to sample a node in this lookout and eventually create a connected roadmap in $F'$. 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 has been to reveal this fact.

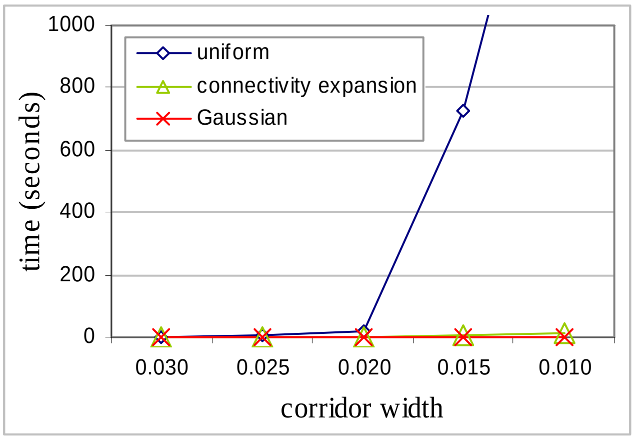


**Fig. 4.** Comparison of three strategies with different sampling measures. The plot shows the average running times over 30 runs on the problem in Fig. 3 as the corridor width decreases.

According to Theorems 1–3, expansiveness provides a sufficient condition for `Basic-PRM` work well. One can also prove that expansiveness is necessary in the following sense: if $F$ is not expansive for large enough values of $\varepsilon$, $\alpha$, and $\beta$, then it is always possible to choose a query in $F$ so that `Basic-PRM` fails to succeed with high probability. This indicates that expansiveness is a good characterization of the complexity of the free space for PRM planning. We do not have a 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 $\varepsilon$-complexity) and they are more specific. However, since the values of $\varepsilon$, $\alpha$, and $\beta$ 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?

In the previous section, we have analyzed the performance of `Basic-PRM` when the uniform sampling measure is used. However, most PRM planners employ non-uniform sampling measures that dramatically improve performance. To illustrate, Fig. 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 [12], and the Gaussian strategy [2]. 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?

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, the reason 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:

- The *two-phase connectivity expansion strategy* [12] 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 Fig. 5*a* around the corridor.
- In each sampling operation, the *Gaussian strategy* [2] 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* (Fig. 5*b*). The rationale is that points inside narrow passages, which have poor visibility, often lie near the boundary. Focusing on the boundary may increase the sampling density inside narrow passages.

Fig. 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.

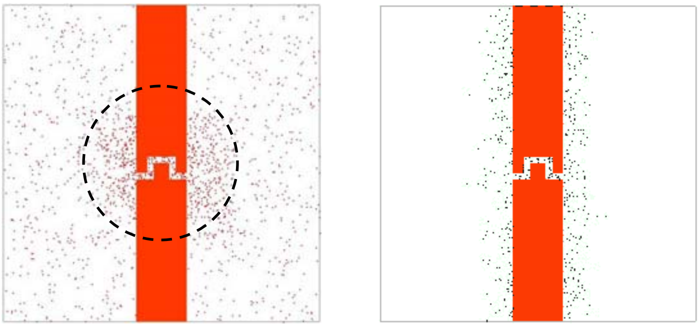

(*a*) (*b*)

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

The above two non-uniform strategies are chosen here only for illustration. Other strategies have been proposed, and some of them achieve even greater speedup. They use various techniques to increase sampling density in subsets of *F* expected to have poor visibility. For instance, the bridge test extends the Gaussian strategy and samples three configurations, instead of two, to better identify narrow passages [7]. Other techniques identify narrow passages in a robot's workspace (*e.g.*, by computing the medial axis) and use this information to sample more densely in regions of *F* likely to contain narrow passages [4, 5, 20]. For a robot manipulator arm, it has been shown that over-sampling near singular configurations improves performance [14]. Indeed, at a singular configuration $\boldsymbol{q}_s$, the arm's end-effector loses some degrees of freedom. Thus the region of *F* near $\boldsymbol{q}_s$ has a flattened shape, resulting in poor visibility. Instead of using heuristics to locate regions with poor visibility, an alternative is to check directly the definition of visibility to prune a roadmap and avoid wasting effort in regions with good visibility [19], but this may involve high computational cost. A quite different approach is to slightly dilate *F* [8, 18]. As visibility in dilated *F* is more favorable, planning becomes easier. A path found in the dilated space is then deformed into one in *F*.

## 5 How important is the sampling source?

We have mentioned in Section 2 that a sampling strategy ($\pi$, S) is characterized by a probability measure $\pi$ 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 sequence of numbers that closely approximate the statistical properties of true random numbers. In particular, a pseudo-random sequence is slightly irregular to

| | $\pi_U$ | $\pi_G$ |
|---|---|---|
| $S_{ran}$ | 1.0 | 92 |
| $S_{hal}$ | 1.4 | 23 |
| $S_{opt}$ | 4.0 | 92 |

(*a*)

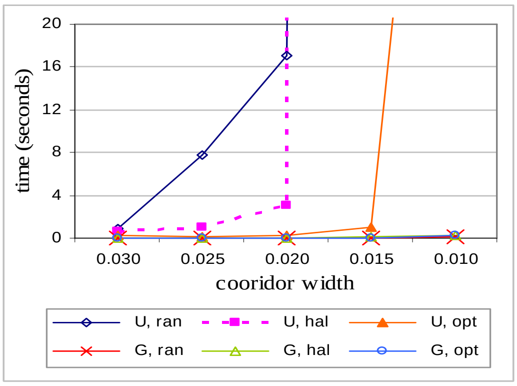


(*b*)

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

simulate the effect that each number is chosen independently. Note that if we fix the seed of a pseudo-random source, the numbers generated are in fact deterministic. To get multiple independent runs of a PRM planner, we must use a different seed for each run. 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 [13]. 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 with that of 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 $S_{hal}$ [17] and the incremental discrepancy-optimal sequence $S_{opt}$ [15], both of which have been reported to often outperform $S_{ran}$ [6, 13, 15]. We then pair each source with two probability measures, the uniform measure $\pi_U$ and the measure $\pi_G$ 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.

- **The sampling measure versus the sampling source.** Fig. 6*a* compares the six strategies on the example in Fig. 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 $(\pi_U,S_{ran})$ versus that of the strategy of the entry. So, the table reports the *speedup* over $(\pi_U,S_{ran})$. The running times for $(\pi_U,S_{ran})$ and $(\pi_G,S_{ran})$ are averaged over 30 independent runs. The second column $(\pi_U)$ shows that $S_{hal}$ and $S_{opt}$ indeed achieve some speedup over $S_{ran}$, but far greater speedup is achieved by switching to $\pi_G$. Furthermore, the advantage of $S_{hal}$ and $S_{opt}$ over $S_{ran}$ observed with $\pi_U$ vanishes when we switch to $\pi_G$. These results are reinforced in Fig. 6*b*, which plots the running times of the six strategies, as the corridor width decreases. The three

| | $\pi_U$ | $\pi_G$ |
|---|---|---|
| $S_{ran}$ | 1.0 | 40.3 |
| $S_{hal}$ | 3.9 | 33.2 |
| $S_{opt}$ | 0.9 | 42.2 |

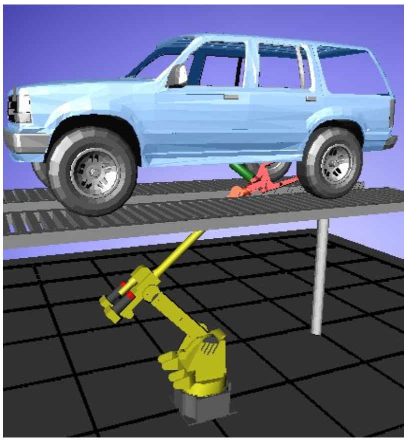

**Fig. 7.** Comparison of six sampling strategies on a more realistic problem.

curves bundled together at the bottom of the plot all correspond to strategies using $\pi_G$, 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 Fig. 7, 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.

- **Dependence on dimensionality.** The main basis for deterministic sources is that they minimize criteria such as discrepancy or dispersion. However, the computational cost of maintaining a fixed discrepancy or dispersion increases *exponentially* with dim($C$) [17]. The samples generated by a deterministic source distribute evenly and regularly over $[0,1]^{\dim(C)}$, and so they roughly correspond to a grid with $N^{1/\dim(C)}$ discretized intervals per axis, where $N$ is the number of samples. In typical PRM planning problems, $N$ is relatively small, while dim($C$) could be large (greater than 6). This leads to large discrepancy and dispersion, even when a deterministic source is used. Hence, the advantage that deterministic sources can possibly achieve over pseudo-random sources fades away as dim($C$) increases. Fig. 8 gives an example, showing the running times of the six strategies as dim($C$) increases from 3 to 8. The robot is a planar linkage with a mobile base. We increase dim($C$) by adding more links Fig. 8 shows that the running time of $(\pi_U,S_{opt})$ increases quickly with dim($C$). The increase is slower with $(\pi_U,S_{hal})$ and even slower with $(\pi_U,S_{ran})$. It is interesting to observe that $(\pi_U,S_{hal})$ performs slightly better than $(\pi_U,S_{ran})$ when $\dim(C) \le 6$, but worsens afterwards (see the inset in the plot). The three strategies using $\pi_G$ all have only moderate increases in running times. As dim($C$) increases, visibility properties become less uniformly favorable over $F$, and the advantage of $\pi_G$ over $\pi_U$ grows.

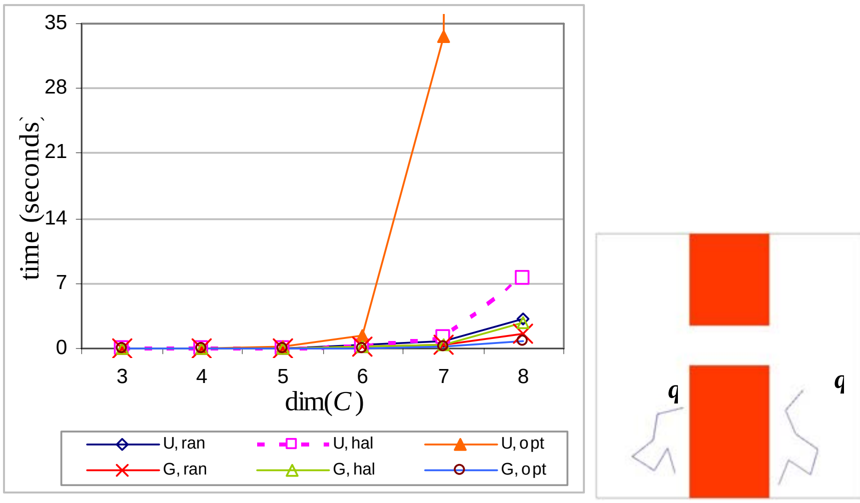


**Fig. 8.** The dependence of six sampling strategies on dim(*C*). The inset in the left plot is a zoom of the lower portions of the curves.

## 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.

To speed up PRM planning, one promising research direction 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 $\pi$ continuously. Initial work along this line has appeared recently [1, 10]. In [1], an approximate model of the configuration space is built and used to sample configurations so that the expected value of a utility function is maximized. A crucial issue here is to define a utility function that closely approximates the expected cost of reaching a good roadmap. In [10], the sampling measure $\pi$ is constructed as a linearly weighted combination of component measures with complementary strengths. To adjust $\pi$, the weights are updated after each sampling operation during roadmap construction to favor the component measures that give the most promising results. An important issue here is then to develop good criteria to assess the performance of component measures.

**Acknowledgements:** D. Hsu is supported by NUS grant R252-000-145-112. J.C. Latombe is supported by NSF grants ACI-0205671, IIS-0412884, and DMS-0443939, and NIH grant 5R33 LM007295. H. Kurniawati is supported an NUS Postgraduate Scholarship. We thank T. Bretl, L. Kavraki, M. Saha, and G. Sánchez-Ante for their comments on early drafts of this paper. We also thank S. Lindemann and S. LaValle for making available the code for generating deterministic low-discrepancy sequences.

## References

1. B. Burns and O. Brock. Sampling-based motion planning using predictive models. In *Proc. IEEE Int. Conf. on Robotics & Automation*, 2005.
2. 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.
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,* Chapter 7, MIT Press, 2005.
4. M. Foskey, M. Garber, 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.
5. 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.
6. 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.
7. 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.
8. D. Hsu, L.E. Kavraki, J.C. Latombe, R. Motwani, and S. Sorkin. On finding narrow passages with probabilistic roadmap planners. In P.K. Agarwal et al., editors, *Robotics: The Algorithmic Perspective: 1998 Workshop on the Algorithmic Foundations of Robotics*, pages 141–154. A. K. Peters, Wellesley, MA, 1998.
9. 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.
10. 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.
11. 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.

12. L.E. Kavraki, P. Švestka, 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.
13. 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.
14. P. Leven and S. Hutchinson. Using manipulability to bias sampling during the construction of probabilistic roadmaps. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 2134–2140, 2002.
15. S.R. Lindemann and S.M. LaValle. Incremental low-discrepancy lattice methods for motion planning. In *Proc. IEEE Int. Conf. on Robotics & Automation*, pages 2920–2927, 2003.
16. M. Lin and D. Manocha. Collision and proximity queries. In J.E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 35. CRC Press, 2004.
17. J. Matousek. *Geometric Discrepancy*. Springer-Verlag, Berlin, 1999.
18. M. Saha and J.C. Latombe. Finding narrow passages with probabilistic roadmaps: The small step retraction method. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2005.
19. T. Siméon, J.P. Laumond, and C. Nissoux. Visibility-based probabilistic roadmaps for motion planning. *J. of Advanced Robotics*, 14(6):477-494, 2000.
20. 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.